

Computational strategies for mapping equilibrium phase diagrams

A.D. Bruce⁽¹⁾ and N.B. Wilding⁽²⁾

⁽¹⁾ *Department of Physics and Astronomy, The University of Edinburgh
Edinburgh, EH9 3JZ, United Kingdom*

⁽²⁾ *Department of Physics, University of Bath
Bath, BA2 7AY, United Kingdom*

We survey the portfolio of computational strategies available for tackling the generic problems of phase behavior – free-energy-estimation and coexistence-curve mapping.

PACS numbers: 05.70.Fh, 64, 02.70, 05.10

I. INTRODUCTION: DEFINING OUR PROBLEM

The phenomenon of *phase behavior* is generic to science; so we should begin by defining it in a general way: it is the organization of many-body systems into forms which reflect the interplay between constraints imposed macroscopically (through the prevailing external conditions) and microscopically (through the interactions between the elementary constituents). In addition to its most familiar manifestations in condensed matter science the phenomenon –and its attendant challenges– features in areas as diverse as gauge-theories of sub-nuclear structure, the folding of proteins and the self-organization of neural networks.

We will of course be rather more focused here. We shall be concerned with the *generic computational strategies* needed to address the problems of phase behavior. The physical context we shall explore will not extend beyond the *structural* organization of the elementary phases (liquid, vapor, crystalline) of matter, although the *strategies* are much more widely applicable than this. We shall have nothing to say about a wide spectrum of techniques (density functional theory [1], integral equation theories [2], anharmonic perturbation theory [3], virial expansions [4]) which have done much to advance our understanding of sub-classes (solid-liquid, solid-solid or liquid-gas) of phase behavior, but which are less than ‘generic’. For the most part we shall also largely restrict ourselves to systems comprising simple classical particles in defect-free structures, interacting through a prescribed potential function.

There is one other rather more significant respect in which our objectives are limited; we need to identify it now. In its fullest sense the ‘problem’ of phase behavior entails questions of *kinetics*: *how* phase transformations occur. Within a computational framework such questions are naturally addressed with the techniques of Molecular Dynamics (MD), which numerically integrate the equations of motion associated with the many-body-potential, and thereby attempt to replicate, authentically, the phase-transformation *process* itself. There is a long, rich and distinguished history of activity of this kind, providing insights into many (perhaps the most) challenging and interesting issues associated with phase behavior. But the ‘authenticity’ carries a price: like their laboratory-counterparts, such computer experiments display phenomena (hysteresis and metastability) associated with the long time-scales required for the reorganization processes. These phenomena are useful qualitative signatures of a phase transformation and worthy of study in their own right. But they tend to obscure the *intrinsically sharp* characteristics of the transformation (Fig. 1).

If our interest is restricted to those sharp characteristics –in particular *where* a phase transformation occurs in an ‘ideal experiment’– we may formulate (and in principle solve) the problem entirely within the framework of the equilibrium statistical mechanics of the competing phases (the framework which implies that ‘sharpness’).

This is the stance we adopt here: we restrict our attention to the task of mapping *equilibrium phase boundaries*. This choice leads naturally to another. Freed of the need to capture the authentic dynamics, one finds that the natural computational tool becomes Monte Carlo (MC) rather than MD. The strategic advantage of this choice is the range of ways (not all of which have yet been dreamt of) in which MC algorithms may be engineered to move around configuration space. Indeed, understanding the distinctive features of that ‘space’ which are implied by the *occurrence* of phase behavior, and engineering an algorithm (effectively a pseudo-dynamics) to match, are the key themes of recent activities, which we shall highlight in this work.

We shall begin (Section II) by assembling the basic equipment. Section II A formulates the problem in the complementary languages of thermodynamics and statistical mechanics. The shift in perspective –from ‘free energies’ in the former to ‘probabilities’ in the latter– helps to show what the *core problem* of phase behavior really is: a comparison of the *a priori* probabilities of two regions of configuration space. Section II B outlines the standard portfolio of MC tools and explains why they are not equal to the challenge posed by this core problem.

Section III identifies and explores what proves to be *the* key concept here – the notion of a *path through configuration space*. Most approaches to the core problem utilize a ‘path’ which, in some sense, ‘connects’ the two regions of interest.

They can be classified, helpfully we think, on the basis of the choices made in regard to that path: the *route* it follows through configuration space; and the way in which one contrives to *sample* the configurations along that route. There are a limited number of generically-distinct path-routing options; we summarize them in Section III B. In Section III C we identify the generically distinct path-sampling strategies. There are actually fewer than one might suppose; this is an area in which many wheels have been reinvented. Some of the new wheels at least run more quickly than their precursors; the advances here (the refinement of ‘extended sampling’ techniques) account for much of the recent activity in this area. They are essential for would-be practitioners but not for the logic of the story. Accordingly they are addressed separately, in an Appendix.

The relatively few generic choices in regard to routing and sampling have been deployed in combinations too numerous to mention let alone discuss. We have, instead, elected to survey a small number of what we shall call *path-based techniques* (combinations of routing and sampling strategies) in some detail. This survey (Section IV) extends from the standard technique of long-standing, numerical integration to reference macrostates, to one only recently introduced, in which the core problem is solved by a MC leap directly between the phases of interest.

We have chosen to organize our strategic survey around the concept of a path. There are some strategies which do not naturally fit into this framework, but which must certainly be included here: we discuss them in Section V.

Solving the core problem at some state point is the major part of the task of determining the phase-boundary; but not quite all. To complete the job one has to *track* a phase boundary (once one has found a point on it) and one has to *extrapolate* from the limited sizes of system that simulations can handle to the thermodynamic limit of interest. Section VI reviews the complementary tools needed.

We have chosen for the most part to focus on relatively idealised model systems; in Section VII we consider some of the issues associated with departures from the ideal.

Finally, section VIII offers some thoughts on how we are likely to make further progress.

There are three further caveats about what the reader can hope to find in what follows. It is not easy (perhaps not possible) to be exhaustive and clear; we aim to be clear. One can adopt an organizational structure founded either on the history or the logic of the ideas; we have chosen the latter. We have surely devoted disproportionately much space to our own contributions; this is mainly because we can explain them best.

II. BASIC EQUIPMENT

A. Formulation: statistical mechanics & thermodynamics

We will formulate the problem simply but generally. Consider a system of structureless, classical particles, characterized macroscopically by a set of thermodynamic coordinates (such as the temperature T) and microscopically by a set of model parameters which prescribe their interactions. The two sets of parameters play a strategically similar role; it is therefore convenient to denote them, collectively, by a single label, c (for ‘conditions or ‘constraints’ or ‘control parameters’ in thermodynamic-and-model space).

We are fundamentally concerned with the phase behavior rooted in the *spatial* organization of the particles and reflected in the statistical behavior of their position coordinates $\vec{r}_i, i = 1, N$. The components of these coordinates are the principal members of a set of *generalized coordinates* $\{q\}$ locating the system in its *configuration space*. In some instances (dealing with fluid phases) it is advantageous to work with ensembles in which particle number N or system volume V is free to fluctuate; the coordinate set $\{q\}$ is then extended accordingly (to include N or V) and the control parameters c extended to include the corresponding fields (chemical potential μ or pressure P).

The statistical behavior of interest is encapsulated in the equilibrium probability density function $P_0(\{q\}|c)$. This PDF is determined by an appropriate ensemble-dependent, dimensionless [6] configurational energy $\mathcal{E}(\{q\}, c)$. The relationship takes the generic form

$$P_0(\{q\}|c) = \frac{1}{Z(c)} e^{-\mathcal{E}(\{q\}, c)} \quad (1)$$

where the normalizing prefactor (the partition function) is defined by

$$Z(c) \equiv \int \prod_i dq_i e^{-\mathcal{E}(\{q\}, c)} \quad (2)$$

Some small print should appear here; we shall come back to it. The different phases which the system displays will in general be distinguished by the values of some macroscopic property loosely described as an *order parameter*. Thus, for example, the density serves to distinguish a liquid from a vapor; a structure factor distinguishes a liquid from a

crystalline solid. A suitable order parameter, M , allows us to associate with each phase, α , a corresponding portion $\{q\}_\alpha$ of $\{q\}$ -space. We write that statement concisely in the form:

$$\{q\} \in \{q\}_\alpha \text{ iff } M(\{q\}) \in [M]_\alpha \quad (3)$$

where $[M]_\alpha$ is the set of order parameter values consistent with phase α . The partitioning of $\{q\}$ -space into distinct regions is the key feature of the core problem.

The equilibrium properties of a particular phase α follow from the conditional counterpart of Eq. 1

$$P_0(\{q\}|\alpha, c) = \begin{cases} \frac{1}{Z_\alpha(c)} e^{-\mathcal{E}(\{q\}, c)} & \{q\} \in \{q\}_\alpha \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

with

$$Z_\alpha(c) \equiv \int \prod_i dq_i \Delta_\alpha[M(\{q\})] e^{-\mathcal{E}(\{q\}, c)} \equiv e^{-\mathcal{F}_\alpha(c)} \quad (5)$$

The last equation defines $\mathcal{F}_\alpha(c)$, the *free energy* of phase α , while

$$\Delta_\alpha[M] \equiv \begin{cases} 1 & M \in [M]_\alpha \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

so that the integral is effectively confined to the set of configurations $\{q\}_\alpha$ associated with phase α .

Since the notation does not always make it apparent, we should note that α and c play operationally similar roles as *macrostate labels*: together they identify the distinct sets of equilibrium macroscopic properties emerging from the equilibrium distribution, Eq. 1. For some purposes we will find it useful to concatenate the two labels into a single ‘grand’ macrostate label

$$\alpha, c \longrightarrow \mathcal{C} \quad (7)$$

For the time being we shall continue to display the two labels separately.

Now appealing back to Eq. 1 we may write

$$Z_\alpha(c) = \int \prod_i dq_i \Delta_\alpha[M] e^{-\mathcal{E}(\{q\}, c)} = Z(c) \int \prod_i dq_i \Delta_\alpha[M] P_0(\{q\}|c)$$

The *a priori* probability of phase α may thus be related to its free energy by

$$P_0(\alpha|c) \equiv \int \prod_i dq_i \Delta_\alpha[M] P_0(\{q\}|c) = \frac{Z_\alpha(c)}{Z(c)} = \frac{e^{-\mathcal{F}_\alpha(c)}}{Z(c)} \quad (8)$$

Alternatively

$$P_0(\alpha|c) \equiv \int \prod_i dq_i \Delta_\alpha[M] P_0(\{q\}|c) = \int dM \Delta_\alpha[M] P_0(M|c) \quad (9)$$

where $P_0(\mathcal{M}|c)$ is the equilibrium distribution of the chosen order parameter. Though seemingly naught but a tautology this representation proves remarkably fruitful

For two phases, α and $\tilde{\alpha}$ say, it then follows that

$$\Delta\mathcal{F}_{\alpha\tilde{\alpha}}(c) \equiv \mathcal{F}_\alpha(c) - \mathcal{F}_{\tilde{\alpha}}(c) = \ln \frac{Z_{\tilde{\alpha}}(c)}{Z_\alpha(c)} = \ln \frac{P_0(\tilde{\alpha}|c)}{P_0(\alpha|c)} = \ln \frac{\int dM \Delta_{\tilde{\alpha}}[M] P_0(M|c)}{\int dM \Delta_\alpha[M] P_0(M|c)} \quad (10)$$

This is a key equation in several respects: it is conceptually helpful; it is cautionary; and it is suggestive, strategically.

At a conceptual level, Eq. 10 provides a helpful link between the languages of thermodynamics and statistical mechanics. According to the familiar mantra of thermodynamics the favored phase will be that of *minimal free energy*; from a statistical mechanics perspective the favored phase is the one of *maximal probability*, given the probability partitioning implied by Eq. 1.

We also then see (the ‘cautionary’ bit) that the thermodynamic mantra presupposes the validity of Eq. 1, and thence of its ‘small print’, which we must now spell out. In general Eq. 1 presupposes ergodicity on the space $\{q\}$.

The framework can thus be trusted to tell us what we will ‘see’ for some given c (the ‘favored phase’) only to the extent that appropriate kinetic pathways exist to allow sampling (ultimately, *comparison*) of the distinct regions of configuration space associated with the different phases. In the context of laboratory experiments on real systems the relevant pathways typically entail the nucleation and growth of droplets of one phase embedded in the other; the associated time scales are long; and Eq 10 will be relevant only if the measurements extend over correspondingly long times. The fact that they frequently do not is signaled in the phenomena of metastability and hysteresis we have already touched on.

Finally, Eq. 10 helps to shape strategic thinking on how to broach the problem computationally.

It reminds us that what is relevant here is the *difference* between free energies of two competing phases and that this free energy difference is a *ratio* of the *a priori* probabilities of the two phases. It implies that the phase boundary may be identified as the locus of points of equal *a priori* probability of the two phases, and that such points are in principle identifiable through the condition that the order parameter distribution will have equal integrated weights (areas) in the two phases. The discussion of the preceding section also suggests that the pathways by which our simulated system passes between the two regions of configuration space associated with the two phases will play a strategically crucial role. While, for the reasons just discussed, the details of those pathways are essential to the *physical applicability* of Eq. 10 they are *irrelevant* to the *values* of the quantities it defines; we are thus free to engineer whatever pathways we may wish.

These considerations lead one naturally to the Monte Carlo toolkit.

B. Tools: elements of Monte Carlo

The Monte Carlo method probably ranks as the most versatile theoretical tool available for the exploration of many-body systems. It has been the subject both of general pedagogical texts [7] and applications-focused reviews [8]. Here we provide only its elements - enough to understand why, if implemented in its most familiar form, it does not deliver what we need, and to hint at the extended framework needed to make it do so.

The MC method generates a sequence (Markov chain) of configurations in $\{q\}$ -space. The procedure can be constructed to ensure that, in the ‘long-enough-term’, configurations will appear in that chain with *any* probability density, $P_S(\{q\})$ (the ‘S’ stands for ‘sampling’) we care to nominate. The key requirement (it is not strictly necessary [9]; and –as we shall see– it is not always sufficient) is that the transitions, from one configuration $\{q\}$ to another $\{q'\}$, should respect the detailed balance condition

$$P_S(\{q\})P_S(\{q\} \rightarrow \{q'\}) = P_S(\{q'\})P_S(\{q'\} \rightarrow \{q\}) \quad (11)$$

where $P_S(\{q\} \rightarrow \{q'\})$ is the transition probability, the probability density of configuration $\{q'\}$ at Markov chain step $t + 1$ given configuration $\{q\}$ at time t . (We have added a subscript to emphasize that its form is circumscribed by the choice of sampling density, through Eq. 11.) MC transitions satisfying this constraint are realized in a two-stage process. In the first stage, one generates a trial configuration $\{q'\} = \mathcal{T}\{q\}$, where \mathcal{T} is some generally stochastic selection procedure; the probability density of a trial configuration $\{q'\}$ given $\{q\}$ is of the form

$$P_{\mathcal{T}}(\{q'\} | \{q\}) = \langle \delta(\{q'\} - \mathcal{T}\{q\}) \rangle_{\mathcal{T}} \quad (12)$$

where $\langle \cdot \rangle_{\mathcal{T}}$ represents an average with respect to the stochastic variables implicit in the procedure \mathcal{T} . In the second stage the ‘trial’ configuration is accepted (the system ‘moves’ from $\{q\}$ to $\{q'\}$ in configuration space) with probability P_A , and is otherwise rejected (so the system ‘stays’ at $\{q\}$); the form of the acceptance probability is prescribed by our choices for P_S and $P_{\mathcal{T}}$ since

$$P_S(\{q\} \rightarrow \{q'\}) = P_{\mathcal{T}}(\{q'\} | \{q\})P_A(\{q\} \rightarrow \{q'\})$$

It is then easy to verify that the detailed balance condition (Eq. 11) is satisfied, if the acceptance probability is chosen as

$$P_A(\{q\} \rightarrow \{q'\}) = \min \left\{ 1, \frac{P_S(\{q'\})P_{\mathcal{T}}(\{q\} | \{q'\})}{P_S(\{q\})P_{\mathcal{T}}(\{q'\} | \{q\})} \right\} \quad (13)$$

Suppose that, in this way, we build a Markov chain comprising a total of t_T steps; we set aside the first t_E configurations visited; we denote by $\{q\}^{(t)}$ ($t = 1, \dots, t_U$) the configurations associated with the subsequent $t_U \equiv t_T - t_E$ steps. The promise on the MC package is that the expectation value $\langle Q \rangle_S$ of some observable $Q = Q(\{q\})$ defined by

$$\langle Q \rangle_S = \int \prod_i dq_i P_S(\{q\}) Q(\{q\}) \quad (14)$$

may be *estimated* by the sample average

$$\langle Q \rangle_S \stackrel{eb}{=} \frac{1}{t_U} \sum_{t=1}^{t_U} Q(\{q\}^{(t)}) \quad (15)$$

Now we must consider the choices of P_S and P_T . Tailoring those choices to whatever task one has in hand provides potentially limitless opportunity for ingenuity. But at this point we consider only the simplest possibilities. The sampling distribution $P_S(\{q\})$ is chosen to be the appropriate *equilibrium* distribution $P_0(\{q\} | c)$ (Eq. 1) so that the configurations visited are representative of a ‘real’ system, even though their sequence is not an authentic representation of the ‘real’ dynamics. We shall refer to this form of sampling distribution as *canonical*. The trial-coordinate selection procedure \mathcal{T} is chosen to comprise some *small* change of *one* coordinate; the change is chosen to be small enough to guarantee a reasonable acceptance probability (Eq. 13), but no smaller, or the Markov chain will wander unnecessarily slowly through the configuration space. We shall refer to this form of selection procedure as *local*. For such schemes (and sometimes for others) the selection probability density typically has the symmetry:

$$P_T(\{q\} \rightarrow \{q'\}) = P_T(\{q'\} \rightarrow \{q\}) \quad (16)$$

With these choices, Eq. 13 becomes

$$P_A(\{q\} \rightarrow \{q'\}) = \mathcal{A}(\Delta\mathcal{E}) \quad (17)$$

where

$$\Delta\mathcal{E} \equiv \mathcal{E}(\{q'\}, c) - \mathcal{E}(\{q\}, c) \quad (18)$$

and

$$\mathcal{A}(x) \equiv \min\{1, \exp[-x]\} \quad (19)$$

defines the Metropolis acceptance function [10].

These choices are not only the simplest, they are also the most frequent: the local-canonical strategy is the staple Monte Carlo method, and has contributed enormously to our knowledge of many-body systems.

From what we have said it would seem that this staple strategy would also deliver what we require here. If, as promised, the Markov chain visits configurations with the canonical probability (Eq. 1) we should merely have to determine

$$P_0(\alpha|c) \equiv \langle \Delta_\alpha \rangle_0 \quad (20)$$

from its estimator (Eq. 15), effectively the proportion of time the system is found in region $\{q\}_\alpha$. Eq. 10 would then take care of the rest. But the local-canonical strategy fails us here. The Markov chain typically does not extend beyond the particular region of configuration space $\{q\}_\alpha$ in which it is initiated and the distribution of the (any) ‘order parameter’ will capture only the contributions associated with that phase. The observations thus provide no basis for assigning a value to the relative probabilities of the two phases, and thus of estimating the location of the phase boundary.

This failure is a reflection of *both* of our ‘simple’ choices. Firstly, the *local* character of coordinate updating yields a dynamics which (though scarcely authentic) shares the essential problematic feature of the kinetic pathways supported by ‘real’ dynamics: evolution from one phase to another will typically require a traverse through intermediate regions in which the configurations have a two-phase character. Secondly, the choice of a *canonical* sampling distribution ensures that the probability of such intermediate configurations is extremely small; inter-phase traverses occur only on correspondingly long time scales. At this point one remembers the small-print (‘in the long term’) that accompanies the MC toolkit. As a result the effective sampling distribution is not the nominal choice, $P(\{q\}|c)$ (Eq. 1) but $P(\{q\}|\alpha, c)$ (Eq. 4), with the phase α determined by our choice of initial configuration. Since the sampling distribution appears only in the acceptance probability (Eq. 13), and then only as a *ratio*, the algorithm does not distinguish between $P(\{q\}|c)$ and $P(\{q\}|\alpha, c)$ –as long as it is trapped in $\{q\}_\alpha$.

We conclude that the local-canonical MC strategy cannot directly deliver the simultaneous comparison between two phases which (Eq. 10 suggests) provides the most efficient resolution of the phase-boundary problem: in most

circumstances this strategy will simply explore a *single* phase. We must now ask whether we can get by with two *separate* (but still local-canonical) ‘single-phase’ simulations, each determining the free-energy (or, equivalently, partition function) of one phase (Eq. 5).

Let us first be clear about the circumstances in which a ‘single phase’ simulation makes sense. The brief and loose answer is: when the time (Markov chain length) $t_{e\alpha}$ typical of escape from phase α is long compared to the time $t_{s\alpha}$ required for effective sampling of the configuration space of that phase. More fully, and a little more formally: when there exists some $t_{s\alpha} < t_{e\alpha}$ such that the configuration set $\{q\}_\alpha$ defined by Eq. 3 is effectively equivalent to that defined by the condition

$$\{q\} \in \{q\}_\alpha \text{ iff } \{q\} \text{ is reachable from } \{q\}_\alpha^0 \text{ within time } t_{s\alpha} \quad (21)$$

In this formulation, the configurations in $\{q\}_\alpha$ are identified as those that may be reached in a simulation of length $t_{s\alpha}$, initiated from some configuration $\{q\}_\alpha^0$ that is associated with phase α but is otherwise arbitrary. The equivalence of Eqs. 3 and 21 is assumed (usually tacitly) in all ‘single phase’ simulations.

Assuming these conditions are fulfilled, our single-phase local-canonical algorithm will allow us to estimate the single-phase canonical expectation value of any ‘observable’ Q defined on the configurations $\{q\}$

$$\langle Q \rangle_{0,\alpha} = \int \prod_i dq_i P_0(\{q\}|\alpha, c) Q(\{q\}) \quad (22)$$

from the average (Eq. 15) over a sample of the canonically-distributed configurations. But the single-phase partition function Z_α is *not* [11] an ‘average over canonically distributed configurations’; rather, it measures the *total effective weight* of the configurations that contribute significantly to such averages. One can no more deduce it from a sample of single phase properties than one can infer the size of an electorate from a sample of their opinions. Thus local-canonical MC fails to deliver the absolute single-phase free energy also.

To determine the relative stability of two phases under conditions c thus requires a MC framework that, in some sense, does *more* than sample the equilibrium configurations appropriate to the (two) c -macrostates. We have seen where there is room for maneuver – in the choices we make in regard to P_S and P_T . The possibilities inherent in the latter are intuitively obvious: better to find ways of bounding or leaping through configuration space than be limited to the shuffle of local-updating. The fact that we have flexibility in regard to the choice of sampling distribution is perhaps less obvious so it is worth recording the simple result which shows us that we do.

Let P_S and $P_{S'}$ be two arbitrary distributions of the coordinates $\{q\}$. Then the expectation values of some arbitrary observable Q with respect to the two distributions are formally related by the identity

$$\langle Q \rangle_{S'} = \int \prod_i dq_i P_{S'}(\{q\}) Q(\{q\}) = \int \prod_i dq_i P_S(\{q\}) Q(\{q\}) \frac{P_{S'}(\{q\})}{P_S(\{q\})} = \langle \frac{P_{S'}}{P_S} Q \rangle_S \quad (23)$$

Thus, in particular, we can –in principle– determine *canonical* expectation values from an ensemble defined by an *arbitrary* sampling distribution through the relationship

$$\langle Q \rangle_0 = \langle \frac{P_0}{P_S} Q \rangle_S \quad (24)$$

We do not *have* to make the choice $P_S = P_0$. The issue of what sampling distribution will be *optimal* was addressed in the earliest days of computer simulation [12]. The answer depends on the observable Q . In general the ‘obvious’ choice $P_S = P_0$, though not strictly optimal, is adequate. But for some observables the choice of a canonical sampling distribution is *so* ‘sub-optimal’ as to be useless [13]. The core problem we face here has a habit of presenting us with such quantities: we have already seen one example (Eq. 20) and we shall see others.

III. PATHS

There are many ways of motivating, constructing and describing the kind of MC sampling strategy we need; the core idea we shall appeal to here to structure our discussion is that of a *path*.

A. Meaning and specification

For our purposes a *path* comprises a sequence of contiguous macrostates, $\mathcal{C}_1, \mathcal{C}_2 \dots \mathcal{C}_\Omega \equiv \{\mathcal{C}\}$ [14]. By ‘contiguous’ we mean that each adjacent pair in the sequence ($\mathcal{C}_j, \mathcal{C}_{j+1}$ say) have some configurations in common (or that a configuration of one lies arbitrarily close to a configuration of the other). A path thus comprises a quasi-continuous band through configuration space.

The physical quantities that distinguish the macrostates from one another will fall into one or other of two categories, which we shall loosely refer to as *fields*, and *macrovariables*. In the former category we include thermodynamic fields, model parameters and, indeed, the conjugate [15] of any ‘macrovariable’. By ‘macrovariable’ [16] we mean any collective property, aggregating contributions from all or large numbers of the constituent particles, free to fluctuate in the chosen ensemble, but in general sharply-prescribed, in accordance with the Central Limit Theorem. Note that we do not restrict ourselves to quantities that feature on the map of thermodynamics, nor to the parameter space of the physical system itself: with simulation tools at our disposal there are limitless varieties of parameters to vary and properties to observe.

B. Generic routes

It may be evident (it should certainly not be surprising) that the extended MC framework needed to solve the phase-equilibrium problem entails exploration of a path that *links* the macrostates of the two competing phases, for the desired physical conditions \mathcal{C} . The generic choices here are distinguished by the way in which the path is *routed* in relation to the key landmark in the configuration space – the two-phase region which separates the macrostates of the two phases, and which confers on them their (at least meta-) stability. Figure 2 depicts four conceptually different possibilities.

First (Fig. 2(a)) the route may comprise two distinct sections, neither encroaching into the two-phase region, and each terminating in a *reference* macrostate. By *reference* macrostate we mean one whose partition function (and thus free energy) is already known – on the basis of exact calculation or previous measurement. The information accessible through MC study of the two sections of such a path has to be combined with the established properties of the reference macrostates to provide the desired link between the two equilibrium macrostates of interest. This is the traditional strategy for addressing the phase-coexistence problem.

In the case of a liquid-gas phase boundary it is possible to choose a path (Fig. 2(b)) that links the macrostates of the two phases *without* passing through the two-phase region: one simply has to sail around the critical point at which the phase-boundary terminates. The dependence on the existence of an adjacent critical point limits the applicability of this strategy.

The next option (it seems to be the only one left after (a) and (b); but see (d)) is a route which negotiates the probability ravine (free-energy barrier) presented by the two-phase region (Fig. 2(c)). The extended sampling tools we discuss in the next section are essential here; with their refinement in recent years this route has become increasingly attractive.

If either of the phases involved is a *solid*, there are *additional* reasons (we shall discuss them later) to avoid the ravine, over and above the low canonical probability of the macrostates that lie there. It is possible to do so. The necessary strategy (recently developed) is depicted in Fig. 2(d). As in (a) the path comprises two segments, each of which lies within a single phase region. But in contrast to (a) the special macrostates to which these segments lead are not of the traditional reference form: the defining characteristic of *these* macrostates is that they should act as the ends of a ‘wormhole’ along which the system may be transported by a collective Monte Carlo move, from one to phase to the other. In this case extended sampling methods are used to locate the wormhole ends.

C. Generic sampling strategies

The task of exploring (sampling) the macrostates which form a path can be accomplished in a number of conceptually different ways; we identify them here in a rudimentary way. We defer to subsequent sections the discussion of specific examples which will show how the information they provide is used to give the result we seek; and their strengths and weaknesses.

Serial sampling

The most obvious way of gathering information about the set of macrostates $\{\mathcal{C}\}$ is to use canonical Boltzmann sampling to explore them one at a time, with a sampling distribution set to

$$P_S(\{q\}) = P_0(\{q\}|\mathcal{C}_j) \quad (j = 1 \dots \Omega) \quad (25)$$

in turn. One must then combine the information generated in these Ω independent simulations. The traditional approach to our problem (‘integration methods’, Section IV A) employs this strategy; its basic merit is that it is simple.

Parallel sampling

Instead of exploring the path-macrostates serially we may choose to explore them in parallel. In this framework we simulate a set of Ω replicas of the physical system, with the j -th member chosen to realize the conditions \mathcal{C}_j . The sampling distribution in the composite configuration space spanned by $\{q\}^{(1)} \dots \{q\}^{(\Omega)}$ is

$$P_S(\{q\}^{(1)} \dots \{q\}^{(\Omega)}) = \prod_{j=1}^{\Omega} P_0(\{q\}^{(j)}|\mathcal{C}_j) \quad (26)$$

This is not simply a way of exploiting the availability of parallel computing architectures to treat Ω tasks at the same time; more significantly it provides a way of breaking out of the straight-jacket of local-update algorithms. The composite ensemble can be updated through interchanges of the coordinate sets associated with adjacent macrostates (j and $j + 1$ say) to give updated coordinates

$$\{q'\}^{(j)} = \{q\}^{(j+1)} \quad \text{and} \quad \{q'\}^{(j+1)} = \{q\}^{(j)} \quad (27)$$

The chosen sampling distribution (Eq. 26) is realized if such configuration-exchanges are accepted with the appropriate probability (Eqs. 13, 17) reflecting the change incurred in the total energy of the composite system. This change is nominally ‘macroscopic’ (scales with the size of the system) and so the acceptance probability will remain workably large only if the parameters of adjacent macrostates are chosen sufficiently close to one another. In practice this means utilizing a number Ω of macrostates that is proportional to \sqrt{N} [17].

Interchanging the configurations of adjacent replicas is one instance (see Section IVD for another) of a *global-update* in which all the coordinates evolve simultaneously. The pay-off, potentially, is a more rapid evolution around coordinate space –more formally a stronger mixing of the Markov chain. This is of course a general desideratum of any MC framework. Thus it is not surprising that algorithms of this kind have been independently devised in a wide variety of disciplines, applied in a correspondingly wide set of contexts ...and given a whole set of different names [17]. These include Metropolis Coupled Chain [18], Exchange Monte Carlo [19], and Parallel Tempering [20]. In Section IV B we shall see how one variant has been applied to deal with the phase-coexistence problem in systems with a critical point.

Extended sampling

We will use the term *extended sampling* (ES) to refer to an algorithm that allows exploration of a region of configuration space which is ‘extended’ with respect to the range spanned by canonical Boltzmann sampling –specifically, one which assembles the statistical properties of our set of macrostates $\mathcal{C}_1 \dots \mathcal{C}_\Omega$ within a *single* simulation. Again it is straightforward to write down the generic form of a sampling distribution that will achieve this end; we need only a ‘superposition’ of the canonical sampling distributions for the set of macrostates.

$$P_S(\{q\}) = W_0 \sum_{j=1}^{\Omega} P_0(\{q\}|\mathcal{C}_j) \quad (28)$$

where W_0 is a normalization constant. The superficial similarity between this form and those prescribed in Eqs 25 and 26 camouflages a crucial difference. Each of the distributions $P_0(\{q\}|\mathcal{C}_j)$ involves a normalization constant identifiable as (Eqs. 1, 5)

$$w_j = Z(\mathcal{C}_j)^{-1} \quad (29)$$

We do not need to know the set of normalization constants $\{w\}$ to implement *serial* sampling (Eq. 25) since each features *only* as the normalization constant for a *sampling* distribution, which the MC framework does not require. Nor do we need these constants in implementing *parallel* sampling (Eq. 26) since in this case they feature (through their product) only in the one overall normalization constant for the sampling distribution. But Eq. 28 is different. Writing it out more explicitly

$$P_S(\{q\}) = W_0 \sum_{j=1}^{\Omega} w_j e^{-\mathcal{E}(\{q\}, \mathcal{C}_j)} \quad (30)$$

we see that the weights $\{w\}$ control the relative contributions which the macrostates make to the sampling distribution. While we are in principle at liberty to choose whatever mixture we please (we do not *have* to make the assignment prescribed by Eq. 29) it should be clear intuitively (we shall develop this point in section IV C 1) that the choice should confer *roughly* equal probabilities on each macrostate, so that all are well-sampled. It is not hard to see that the weight-assignment made in Eq. 29 is in fact what we need to fulfil this requirement. Evidently, to make *extended* sampling work we *do* need to ‘know’ the weights $w_j = Z(\mathcal{C}_j)^{-1}$. There is an element of circularity here which needs to be recognized. Our prime objective is to determine the (relative) configurational weights of *two* macrostates (those associated with two different phases, under the same physical conditions [21]); to do so (somehow or other –we haven’t yet said how) by extended sampling requires knowledge of the configurational weights of a *whole path’s-worth* of macrostates. There *is* progress here nevertheless. While the two macrostates of interest are remote from one another, the path (by construction) comprises macrostates which are contiguous; it is relatively easy to determine the relative weights of pairs of contiguous macrostates, and thence the relative weights of all in the set. In effect the extended sampling framework allows us to replace one hard problem with a large number of somewhat easier problems.

The machinery needed to turn this general strategy into a practical method (‘building the extended sampling distribution’ or ‘determining the macrostate weights’) has evolved over the years from a process of trial and error to algorithms that are systematic and to some extent self-monitoring. The workings of the machinery is more interesting than it might sound; we will discuss some aspects of what is involved in Sections IV C and IV D. But we relegate more technical discussion (focused on recent advances) to Appendix A. Here we continue with a broader brush.

It would be hard to write a definitive account of the development of extended sampling methods; we will not attempt to do so. The seminal ideas are probably correctly attributed to Torrie and Valleau [22] who coined the terminology *umbrella sampling*. The huge literature of subsequent advances and rediscoveries may be rationalized a little by dividing it into two, according to how the macrostates to be weighted are defined.

If the macrostates are defined by a set of values $[\lambda]$ of some generalized ‘field’ λ , the sampling distribution is of the form

$$P_S(\{q\}) = W_0 \sum_{j=1}^{\Omega} w_j e^{-\mathcal{E}(\{q\}, \lambda_j)} \quad (31)$$

Extended sampling strategies utilizing this kind of representation feature in the literature with a variety of titles: *expanded ensemble* [23], *simulated tempering* [24], *temperature scaling* [25].

On the other hand, if the macrostates are defined on some ‘macrovariable’ M the sampling distribution is of the form

$$P_S(\{q\}) = W_0 \sum_{j=1}^{\Omega} w_j e^{-\mathcal{E}(\{q\})} \Delta_j[M] \quad (32)$$

where

$$\Delta_j[M] \equiv \begin{cases} 1 & M \in \text{range associated with } \mathcal{C}_j \\ 0 & \text{otherwise} \end{cases} \quad (33)$$

Realizations of this formalism go under the names *adaptive umbrella sampling* [26] and the *multicanonical ensemble* introduced by Berg and co-workers [27]. It seems right to attribute the recent revival in interest in extended sampling to the latter work.

In Sections IV C and IV D we shall see that extended sampling strategies provide a rich variety of ways of tackling the phase coexistence problem, including the distinctive problems arising when one of the phases is of solid form.

IV. PATH-BASED TECHNIQUES: A GUIDED TOUR

We now proceed to explore how the strategic options in regard to routing and sampling of paths (Sections IIIB and IIIC) can be fused together to form practical techniques for addressing the phase-coexistence problem. We do not set out to be exhaustive here: there are very many pairings of routes and sampling strategies (the choices are to some extent mutually independent, which is why we discussed them separately). We focus rather on a few key cases; we explain what information is gathered by the chosen sampling of the chosen path and how it is used to yield the desired free-energy comparison; and we assess the strengths and weaknesses of each technique as we go.

A. Keeping it simple: numerical integration and reference states

1. The strategy

The staple approach to the phase-coexistence problem involves serial sampling (Section IIIC) along a path of the type depicted in Figure 2(a). Effectively the single core problem (comparing the configurational weights –free-energies– of the two physical macrostates) is split into two, each requiring comparison of a physical macrostate (c, α) with some suitable reference macrostate, $\mathcal{C}_\alpha^{\text{REF}}$. A reference macrostate will be ‘suitable’ if one can identify a path parameterized by some field λ linking it to the physical macrostate, with $\lambda = \lambda_1$ and $\lambda = \lambda_\Omega$ denoting respectively the physical and the reference macrostates. The sampling distribution at some arbitrary point, λ , on this path is then of the form (Eqs. 4,5, 25)

$$P_S(\{q\}) = P_0(\{q\}|\lambda) = \frac{1}{Z_\alpha(\lambda)} e^{-\mathcal{E}(\{q\}, \lambda)} \quad (34)$$

with

$$Z_\alpha(\lambda) \equiv \int \prod_i dq_i \Delta_\alpha[M(\{q\})] e^{-\mathcal{E}(\{q\}, \lambda)} \equiv e^{-\mathcal{F}_\alpha(\lambda)} \quad (35)$$

We have seen that free-energies like $\mathcal{F}_\alpha(\lambda)$ are not themselves naturally expressible as canonical averages; but their *derivatives* with respect to field-parameters *are* expressible this way. Specifically,

$$\begin{aligned} \frac{\partial \mathcal{F}_\alpha(\lambda)}{\partial \lambda} &= -\frac{1}{Z_\alpha(\lambda)} \frac{\partial Z_\alpha(\lambda)}{\partial \lambda} \\ &= \frac{1}{Z_\alpha(\lambda)} \int \prod_i dq_i \Delta_\alpha[M(\{q\})] \frac{\partial \mathcal{E}(\{q\}, \lambda)}{\partial \lambda} e^{-\mathcal{E}(\{q\}, \lambda)} \\ &= \left\langle \frac{\partial \mathcal{E}(\{q\}, \lambda)}{\partial \lambda} \right\rangle_{\alpha, \lambda} \end{aligned} \quad (36)$$

where the average is to be taken with respect to the canonical distribution for macrostate α, λ (Eq. 34.) The free-energy difference between the physical and reference macrostates is thus formally given by

$$\mathcal{F}(\mathcal{C}_\alpha^{\text{REF}}) - \mathcal{F}_\alpha(c) = \int_{\lambda_1}^{\lambda_\Omega} d\lambda \frac{\partial \mathcal{F}_\alpha(\lambda)}{\partial \lambda} = \int_{\lambda_1}^{\lambda_\Omega} d\lambda \left\langle \frac{\partial \mathcal{E}(\{q\}, \lambda)}{\partial \lambda} \right\rangle_{\alpha, \lambda} \quad (37)$$

A sequence of independent simulations, conducted at a set of points $[\lambda]$ spanning a path from λ_1 to λ_Ω , then allows one to estimate the free energy difference by numerical quadrature:

$$\mathcal{F}(\mathcal{C}_\alpha^{\text{REF}}) - \mathcal{F}_\alpha(c) \stackrel{eb}{=} \sum_{j=1}^{\Omega} \left\langle \frac{\partial \mathcal{E}(\{q\}, \lambda)}{\partial \lambda_j} \right\rangle_{\alpha, \lambda_j} \Delta \lambda \quad (38)$$

This takes us half way. The entire procedure has to be repeated for the second phase $\tilde{\alpha}$, integrating along some path parameterized (in general) by some *other* field $\tilde{\lambda}$ running between the macrostate $\tilde{\alpha}, c$ and some *other* reference state $\mathcal{C}_{\tilde{\alpha}}^{\text{REF}}$. Finally the results of the two procedures are combined to give the quantity of interest (Eq. 10)

$$\Delta\mathcal{F}_{\alpha\tilde{\alpha}}(c) = \mathcal{F}(\mathcal{C}_{\alpha}^{\text{REF}}) - \mathcal{F}(\mathcal{C}_{\tilde{\alpha}}^{\text{REF}}) - \Delta\lambda \sum_{j=1}^{\Omega} \left\langle \frac{\partial \mathcal{E}(\{q\}, \lambda)}{\partial \lambda_j} \right\rangle_{\alpha, \lambda_j} + \Delta\tilde{\lambda} \sum_{j=1}^{\tilde{\Omega}} \left\langle \frac{\partial \mathcal{E}(\{q\}, \tilde{\lambda})}{\partial \tilde{\lambda}_j} \right\rangle_{\tilde{\alpha}, \tilde{\lambda}_j} \quad (39)$$

We shall refer to this strategy as *numerical integration to reference macrostates* (NIRM). It is helpful to assess its strengths and weaknesses armed with an explicit example.

We shall consider what is arguably the archetypal example of the NIRM strategy: the Einstein Solid method (ESM) [28]. The ESM provides a simple way of computing the free energies of crystalline phases, and thence addressing questions of the relative stability of competing crystalline structures. We describe its implementation for the simplest case where the inter-particle interaction is of hard-sphere form; it is readily extended to deal with particles interacting through soft potentials [29].

The name of the method reflects the choice of reference macrostate: a crystalline solid comprising particles which do not interact with one another, but which are bound by harmonic springs to the sites of a crystalline lattice, $\{\vec{R}\}_{\alpha}$, coinciding with that of the phase of interest.

The relevant path is constructed from sampling distributions of the general form prescribed in Eq. 34 with

$$\mathcal{E}(\{q\}, \lambda) = \mathcal{E}(\{\vec{r}\}, c) + \lambda \sum_{i=1}^N \left(\vec{r}_i - \vec{R}_i^{\alpha} \right)^2 \quad (40)$$

The first term on the RHS contains the hard-sphere interactions. The second term embodies the harmonic spring energy, which reflects the displacements of the particles from their lattice sites. In this case the point $\lambda_1 = 0$ locates the physical macrostate. With increasing λ , the energy cost associated with a given set of displacements increases, with a concomitant reduction in the size of the typical displacements. On further increasing λ , one ultimately reaches a point λ_{Ω} beyond which particles are so tightly bound to their lattice sites they practically never collide with one another; the hard sphere interaction term then plays no role, thus realizing the desired reference macrostate, whose free-energy may be computed exactly.

Some of the results of ESM studies of crystalline phases of hard spheres are shown in Table I. We shall discuss them below.

2. Critique

There is much to commend NIRM: it is conceptually simple; it can be implemented with only a modest extension of the simulation framework already needed for standard MC sampling; and it is versatile. It has been applied successfully in free-energy measurements of (*inter alia*) crystalline solids [28], liquids [30] and liquid crystals [31]. It is probably regarded as *the* standard method for attacking the phase-coexistence problem, and it provides the benchmark against which other approaches must be assessed. Nevertheless (as one might guess from the persistence of attempts to develop ‘other approaches’) it is less than ideal in a number of respects. We discuss them in turn.

First, the NIRM method hinges on the *identification of a good path and reference macrostate*. A ‘good’ path is short; but the reference macrostate (the choice of which is limited) may lie far from the physical macrostate of interest, entailing a large number of independent simulations to make the necessary link. In a sense the ESM provides a case in point: the reference macrostate is strictly located at $\lambda_{\Omega} = \infty$; corrections for the use of a finite λ_{Ω} need to be made [28]. This kind of problem is a nuisance, but no worse. A potentially more serious constraint on the path is that the derivative being measured should vary slowly, smoothly and reversibly along it; if it does not the numerical quadrature may be compromised. A phase transition en route (whether in the ‘real’ space of the physical system or the extended-model-space into which NIRM simulations frequently extend) is thus a particular hazard. The realization of NIRM known as the Single Occupancy Cell method (SOCM) [32] provides an example where such concerns arise, and seem not to have been wholly dispelled [33].

The *choice of simulation parameters* also raises issues. Evidently one has to decide how many simulations are to be performed along the path and at which values of λ . In so doing one must strike a suitable balance between minimizing computation time while still ensuring that no region of the path (particularly one in which the integrand varies strongly) is neglected. This may necessitate a degree of trial and error.

The *uncertainties* to be attached to NIRM estimates are problematic in a number of respects. Use of simple numerical quadrature (estimating the integral in Eq. 37 by the sum in Eq. 38) will result in errors. One can reduce such errors by interpolation into the regions of λ between the chosen simulation values (for example using histogram re-weighting techniques discussed in Appendix B); but there will still be systematic errors associated both with the interpolation and with finite sampling times. No reliable and comprehensive prescription for estimating the magnitude of such errors has yet been developed.

These problems are exacerbated when one addresses the quantity of real interest –the *difference* between the free energies of two competing phases. The fact that NIRM treats this as *two* problems rather than *one* (Eq. 39) is problematic in two respects.

First this aspect of the NIRM strategy *compounds* a problem inherent in all simulation studies of this problem (or, indeed, any other in many-body physics): finite-size effects. This issue deserves a section to itself, and gets one (Section VI C). Here we note simply that such effects are harder to assess when one has to synthesize calculations on different phases, utilizing different reference states, and different system sizes –and sometimes conducted by different authors.

Second, since the entire enterprise is constructed so as to locate points (of phase equilibrium) at which the free-energy difference *vanishes*, in NIRM one is inevitably faced with the task of determining some very small number by taking the difference between two relatively large numbers. This point is made more explicitly by the hard-sphere data in Table I. One sees that the *difference* between the values of the free energy [37] of the two crystalline phases is some *four orders of magnitude* smaller than the separate results for the two phases, determined by ESM. Of course one can see this as a testimony to the remarkable care with which the most recent recent ESM studies have been carried out [34]. Alternatively one may see it as a strong indicator that another approach is called for.

B. Parallel tempering; around the critical point

1. The strategy

When the coexistence line of interest terminates in a critical point the two phases can be linked by a single continuous path (Figure 2(b)) which loops around the critical point, eliminating the need for reference macrostates, while still avoiding the inter-phase region. In principle it is possible to establish the location of such a coexistence curve by integration along this route. But the techniques of parallel sampling (Section III C) provide a substantially more elegant way of exploiting such a path, in a technique known as (hyper) parallel tempering (HPT) [38].

Studies of liquid-vapor coexistence are, generally, best addressed in the framework of an open ensemble; thus the state variables here comprise both the particle coordinates $\{\vec{r}\}$ and the particle number N . A path with the appropriate credentials can be constructed by identifying pairs of values of the chemical potential μ and the temperature T which trace out some rough *approximation* to the coexistence curve in the $\mu - T$ plane, but extend into the one-phase region beyond the critical point. Once again there is some circularity here to which we shall return. Making the relevant variables explicit, the sampling distribution (Eq. 26) takes the form

$$P_S(\{\vec{r}\}^{(1)}, N^{(1)} \dots \{\vec{r}\}^{(\Omega)}, N^{(\Omega)}) = \prod_{j=1}^{\Omega} P_0(\{\vec{r}\}^{(j)}, N^{(j)} | \mu_j, T_j) \quad (41)$$

In the context of liquid-vapor coexistence the particle number N (or equivalently the number density $\rho = N/V$) plays the role of an order parameter. Estimates of the distribution $P_0(N | \mu, T)$ are available from the simulation for all the points chosen to define the path. One may then identify the free energy difference from the integrated areas of the branches of this distribution associated with each phase and proceed to search for coexistence using the criterion that these integrated areas should be equal (Eq. 10 et seq.). Figures 3 and 4 show some explicit results [38] for a Lennard-Jones fluid.

2. Critique

The phase-diagram shown in Figure 4 is of course the ultimate objective of such studies; but it is the *distribution* shown in Figure 3 that merits most immediate comment, lest its key feature be taken for-granted. Its ‘key feature’ (which we shall come to recognize as the signature of success in this enterprise) is that it captures the contributions from *both* phases (the two distinct peaks) within the framework of a *single simulation*. The fact that both phases are well-visited shows that this strategy manages to break the ergodic block which dooms conventional MC sampling to explore only the particular phase in which the simulation happens to be launched (Section II B).

The HPT method deals with the ergodic block by avoiding it. The configuration exchange between adjacent replicas (Eq. 27) fuels a form of continuous tempering: a liquid phase configuration resident in a replica low down the coexistence curve may diffuse ‘along the path’ and thus through the replicas, to a point (perhaps super-critical) at which on-going local updating is effective in eroding memory of its liquid origins; that (ever evolving) configuration may then diffuse downwards, to appear in the original replica as a vapor phase configuration.

There is an elegant and powerful idea here, albeit one whose applicability (to the phase-coexistence problem) is limited to systems with critical points. But there is one respect in which it is less than satisfactory –the element of circularity already noted: the method will tell us *if* we have selected a point sufficiently close to coexistence for both phases to have observable probability; it does not in itself tell us what to do if our selection does not satisfy this criterion. How ‘close’ we need to be depends sensitively on the size of the system simulated. The ratio of the two phase probabilities (at a chosen point in field-space) varies exponentially fast with the system size [39]. Thus for a ‘large’ system, unless we are ‘very close’ to coexistence, the equilibrium (grand-canonical) PDF determined in HPT will show signs of only one phase *even when there is no ergodic block preventing access to the other phase*. If the initial choice is close enough to coexistence to provide at least *some* signature of the sub-dominant phase [40] then histogram re-weighting techniques (Appendix B) can be used to give a better estimate of coexistence. But HPT provides no way of systematically improving bad initial estimates, because it provides no mechanism for dealing with the huge difference between the statistical weights of the two phases away from the immediate vicinity of coexistence. To address *that* issue one must turn to extended sampling techniques.

C. Extended sampling: traversing the barrier

1. The strategy

Viewed from the perspectives of configuration space provided by the caricature in Figure 2 the most direct approach to the phase-coexistence problem calls for a full frontal assault on the ergodic barrier that separates the two phases. The extended sampling strategies discussed in Section III C make that possible. The framework we need is a synthesis of Eqs. 10 and 32. We will refer to it generically as Extended Sampling Interface Traverse (ESIT).

Equation 10 shows that we can always accomplish our objective if we can measure the full canonical distribution of an appropriate order parameter. By ‘full’ we mean that the contributions of both phases must be established *and calibrated on the same scale*. Of course it is the last bit that is the problem. (It is always straightforward to determine the two *separately* normalized distributions associated with the two phases, by conventional sampling in each phase in turn.) The reason that it is a problem is that the ‘full canonical’ distribution of the (an) ‘order parameter’ is typically vanishingly small at values intermediate between those characteristic of the two individual phases. The vanishingly small values provide a real, even quantitative, measure of the ergodic barrier between the phases. If the ‘full’ order parameter distribution is to be determined by a ‘direct’ approach (as distinct from the circuitous approach of Section IV B, or the ‘off the map’ approach to be discussed in Section IV D) these low-probability macrostates *must* be visited.

Equation 32 shows how. We need to build a sampling distribution that ‘extends’ along the path of M -macrostates running between the two phases. To do its job that sampling distribution must (Section III C) assign ‘roughly equal’ values to the probabilities of the different macrostates. More explicitly the resulting measured distribution of M -values (following [27] we shall call it multicanonical)

$$P_S(M_j) \equiv \int \prod_i dq_i P_S(\{q\}) \Delta_j[M(\{q\})] \quad (42)$$

should be ‘roughly flat’. It needs to be ‘roughly flat’ because the macrostate of lowest probability sets the size of the bottleneck through which inter-phase traverses must pass. It needs to be *no better* than ‘roughly’ flat because of the way in which (ultimately) it is *used*. It is used to estimate the true canonical distribution $P_0(M)$. The two distributions are simply related by

$$P_0(M_j) \doteq w_j^{-1} P_S(M_j) \quad (43)$$

where $\{w\}$ are the multicanonical weights that define the chosen sampling distribution (Eq. 32) and \doteq means equality to within an overall normalization constant. The procedure by which one uses this equation to estimate the canonical distribution from the measured distribution is variously referred to as ‘unfolding the weights’ or ‘re-weighting’; it is simply one realization of the identity given in Eq. 24. The procedure eliminates any *explicit* dependence on the weights (hence the looseness of the criteria by which they are specified); but it leaves the desired legacy: the relative sizes of the two branches of the *canonical* distribution are determined with a statistical quality that reflects the number of inter-phase traverses in the *multicanonical* ensemble.

This strategy has been applied to the study of a range of coexistence problems, initially focused on lattice models in magnetism [41] and particle physics [42]. Figure 5 [43–45] shows the results of an application to liquid-vapor coexistence in a Lennard-Jones system with the particle number density chosen as an order parameter.

The ESIT strategy has clear advantages with respect to the NIRM and (to a lesser extent) HPT strategies discussed in preceding sections

While HPT *requires* the existence of a critical point ESIT does not (although its existence can be usefully *exploited* to assist in the task of weight generation [44]). Moreover, in contrast to HPT, ESIT does not *require* bootstrapping by a rather good guess of a point on the phase boundary. Extended sampling methods can take huge probability differentials in their stride (100 orders of magnitude is not uncommon) as one can see from the low probability of the macrostates in the inter-phase ravine in Fig. 5. Thus ESIT allows one to measure the ‘free energy difference’ between two phases relatively far from the phase boundary – in fact in any regime in which both phases are at least ‘metastable’ [46].

ESIT has one further potential advantage with respect to HPT, which emerges once one appreciates what it entails at a *microscopic* level –in particular the nature of the intermediate macrostates that lurk in the ravine between the two peaks of the canonical distribution. To do so one needs first to understand a *general* point about the weight generation procedures used to build the sampling distribution. These procedures do not *require* explicit specification of the configurations in the macrostates along the path; rather they *search* for the configurations that dominate those macrostates. The results of the ‘search’ can sometimes be a little unexpected [47,48]. But in the case of the liquid-vapor inter-phase path picked out by choosing the density as order parameter, the results of the search are clear *a priori*: the configurations dominating such a path will show two physically-distinct regions, each housing one of the two phases, separated by an interface. The dominance of this kind of configuration is recognized in arguments which establish the convexity of the free-energy in the thermodynamic limit, and the way in which that limit is approached [49]; this picture also allows one to understand the depth of the probability ravine to be traversed [50]. It follows then that ESIT provides incidental access to these configurations, and thus to information about their dominant feature—the interface. Information of this kind has been exploited in a number of studies [51].

The broader and more far-reaching comparison to be made here is however between NIRM and strategies *like* ESIT which furnish canonical distributions spanning two phases, such as that shown in Figure 5 [52]. Here it seems to us that ESIT wins in two respects. First, it is rather more transparent in regard to uncertainties (in free-energy differences). The error bounds emerging from ESIT (and family) represent purely statistical uncertainties associated with the measurement of the relative weights of two distribution-peaks. In contrast NIRM error bounds have to aggregate the uncertainties (statistical and systematic) associated with different stages of the integration process. Second, it seems rather more satisfying to read-off the result for a free energy difference directly from the likes of Fig. 5 than from a pair of numbers established by appeal to physically-irrelevant reference states.

Nevertheless the ESIT strategy is demanding in a number of respects. Generating the weights is a computationally-intensive job, which is not yet fully self-managing. Subsequent sampling of the resulting multicanonical distribution is a slow process: the dynamics in M -space is a random walk in which visits to the macrostates of low equilibrium probability are secured only at the expense of repeated refusal of proffered moves to the dominant equilibrium macrostates. It helps (though it goes against the spirit of ‘one simulation’) to break the space up into sections, whose length is chosen to reflect an interplay of the diffusion time *between* macrostates and the time associated with relaxation *within* macrostates [53].

These two reservations apply generically to ES methods; the following one is specific to ESIT itself. There are some phase coexistence problems in which an inter-phase path involving an *interface* is computationally fraught. In particular if one of the phases is crystalline (as in the case of melting/freezing) or if both are crystalline (there is a potential *structural* phase transition) such a traverse will involve substantial, physically slow, restructuring — vulnerable to further ergodic traps, and compounding the intrinsic slowness of the multicanonical sampling process [54,55]. In such circumstances it would be better if the inter-phase trip could be accomplished without encountering interfaces. This is possible.

D. From paths to wormholes: phase switch

We shall devote rather more time to this fourth and final example of path-based strategies. We do so for two reasons. First it provides us with an opportunity to touch on a variety of other strands of thought about the ‘free-energy-estimation problem’ which should feature somewhere in this article, and can do so helpfully here. And second, we like it.

The strategy we shall discuss is a way of realizing the direct leap between phases represented in Fig. 2(d). In the literature we have referred to it as ‘Lattice Switch Monte Carlo’ [56,48,57] in the context of phase equilibrium between

crystalline structures, and ‘Phase Switch Monte Carlo’ [58] in the context of solid-liquid coexistence. Here we shall develop the ideas in a general form; and we shall refer to the method as Extended Sampling Phase Switch (ESPS).

1. Strategy

Let us return to the core problem: the evaluation of the *ratio* (Eq. 10) of configurational integrals of the form prescribed in Eq. 5. We can make the problem look different (and possibly make it easier) by choosing to express it in terms of coordinates that are in some sense ‘matched’ to each phase. The simplest useful possibility is provided by an appropriate *linear* transformation

$$\{q\} = \{q\}_\alpha^{\text{REF}} + \tau_\alpha(\{u\}) \quad (44)$$

One can think of $\{q\}_\alpha^{\text{REF}}$ as some reference point in the configuration space of phase α and $\{u\}$ as a ‘displacement’ from that point, modulo some rotation or dilation, prescribed by the operation τ_α . The single-phase partition function in Eq. 5 can then be written in the form

$$Z_\alpha(c) = \det \tau_\alpha \int \prod_i du_i e^{-\mathcal{E}_\alpha(\{u\}, c)} \quad (45)$$

where $\det \tau_\alpha$ is the Jacobean [59] of the transformation (a configuration-independent constant, given the presumed linearity) and

$$e^{-\mathcal{E}_\alpha(\{u\}, c)} = e^{-\mathcal{E}(\{q\}, c)} \Delta_\alpha[\{q\}] \quad (46)$$

The form of the new energy function (‘hamiltonian’) $\mathcal{E}_\alpha(\{u\}, c)$ reflects the *representation* chosen for the region of configuration space relevant to the phase; because this energy function carries a phase label, it can also be used to absorb the constraint (Eq 6) that restricts the integral to that region [60].

The difference between the free energies of the two phases (Eq. 10) can now be written in the form

$$\Delta \mathcal{F}_{\alpha\tilde{\alpha}}(c) = \Delta \mathcal{E}_{\alpha\tilde{\alpha}}^0 - \ln \det \mathbf{S}_{\alpha\tilde{\alpha}} - \ln \mathcal{R}_{\alpha\tilde{\alpha}} \quad (47)$$

where

$$\Delta \mathcal{E}_{\alpha\tilde{\alpha}}^0 = \mathcal{E}(\{q\}_\alpha^{\text{REF}}, c) - \mathcal{E}(\{q\}_{\tilde{\alpha}}^{\text{REF}}, c) \quad (48)$$

is the difference between the hamiltonians of the reference configurations while

$$\mathbf{S}_{\alpha\tilde{\alpha}} = \tau_\alpha \times \tau_{\tilde{\alpha}}^{-1} \quad (49)$$

These two contributions to Eq. 47 are computationally trivial; the computational challenge is now in the third term defined by

$$\mathcal{R}_{\alpha\tilde{\alpha}} = \frac{\int \prod_i du_i e^{-\mathcal{E}_\alpha(\{u\}, c)}}{\int \prod_i du_i e^{-\mathcal{E}_{\tilde{\alpha}}(\{u\}, c)}} \quad (50)$$

In the formulation of Eq. 10 we are confronted with a ratio of configurational integrals defined through the *same* energy function (alias: ‘cost function’ or ‘hamiltonian’) acting on *two explicitly different* regions of configuration space. In contrast Eq. 50 features integrals defined through *different* energy functions acting on *one common* configuration space.

Configurational-integral ratios of the form of Eq. 50 appear widely in the free-energy literature; but the underlying physical motivation is not always the same. The spectrum of possible usages is covered by writing Eq. 50 in the more general form

$$\mathcal{R}_{AB} = \frac{\int \prod_i du_i e^{-\mathcal{E}_A(\{u\})}}{\int \prod_i du_i e^{-\mathcal{E}_B(\{u\})}} \equiv \frac{\mathcal{Z}_A}{\mathcal{Z}_B} \quad (51)$$

where A and B are two generalized macrostate labels, which identify two energy functions.

One meets this kind of ratio (perhaps most naturally) if one considers the difference between the free energies of two macrostates of *one* phase, corresponding to *different* choices of c – as in the influential work of Bennett [61]. And

one meets it if one considers the difference between the free energies of a given model and some approximation to that model, as in the perturbation approach of Zwanzig [62]. Rahman and Jacucci [63] seem to have been amongst the first to consider this structure of problem with the kind of motivation we have given it here –that is, as a way of computing free energy differences between two *phases* directly.

To set the remainder of the discussion in as wide a context as we can we shall develop it in the general notation of Eq. 51, reverting to the specific interpretation of the macrostate labels of primary interest here ($A \rightarrow \alpha, c$; $B \rightarrow \tilde{\alpha}, c$) as appropriate.

In common with most who have addressed the problem posed by this kind of configurational integral ratio, we shall exploit the statistics of an appropriate ‘order-parameter’ defined on the common configuration space and of the general form

$$\mathcal{M}_{AB} \equiv M_A(\{u\}) - M_B(\{u\}) \quad (52)$$

There is considerable license in the choice of M_A and M_B . In the simplest cases it suffices to choose the two energy functions themselves. We make that choice explicitly here so as to expose connections with the work of others. Then the ‘order parameter’ assumes the form

$$\mathcal{M}_{AB} = \mathcal{E}_A(\{u\}) - \mathcal{E}_B(\{u\}) \quad (53)$$

This quantity has the credentials of an ‘order parameter’ in the sense that its behavior is qualitatively different in the two ensembles. To understand this, suppose that we sample from the canonical distribution of the B -ensemble, prescribed by the partition function \mathcal{Z}_B . A ‘typical point’ $\{u\}$ will then characterize a typical configuration of B , in which –for example– no particle penetrates the core-region of the potential of another; that same point $\{u\}$ will, however, describe a configuration of A which is *not* guaranteed to be typical of that ensemble, and *will* in general feature energy-costly regions of core penetration. The order parameter \mathcal{M}_{AB} will thus generally be *positive* in ensemble B ; by the same token it will be *negative* in ensemble A . (We shall see this explicitly in the examples that follow)

One may measure and utilize the statistics of \mathcal{M} in three strategically-different ways.

First, Eqs. 50 and 53 lead immediately to the familiar Zwanzig formula [62],

$$\mathcal{R}_{AB} = \langle e^{-\mathcal{M}_{AB}} \rangle_B = \langle e^{-[\mathcal{E}_A(\{u\}) - \mathcal{E}_B(\{u\})]} \rangle_B \quad (54)$$

In principle then *one single-ensemble-average* suffices to determine the desired ratio. However, this strategy (if unsupported by others) requires [63] that the two ensembles overlap in the sense that, loosely [64], the ‘dominant’ configurations in the A ensemble are a subset of those in the B ensemble. This is a strong constraint; it will seldom if ever be satisfied [65].

The second generic strategy [61] utilizes *two single-ensemble-averages*, that is averages with respect to the *separate* ensembles defined by \mathcal{Z}_A and \mathcal{Z}_B . In particular one may, in principle, measure the canonical probability distributions of the order parameter in each ensemble separately, and exploit the relationship between them [63]

$$\mathcal{R}_{AB} P(\mathcal{M}_{AB} | A) = e^{-\mathcal{M}_{AB}} P(\mathcal{M}_{AB} | B) \quad (55)$$

again presupposing the choice prescribed in Eq. 53. One can re-express this result in an alternative form

$$\mathcal{R}_{AB} = \frac{\langle \mathcal{A}(\mathcal{M}_{AB}) \rangle_B}{\langle \mathcal{A}(\mathcal{M}_{BA}) \rangle_A} \quad (56)$$

where \mathcal{A} is the Metropolis function defined in Eq. 19. Recalling the significance of that function one can see each of the two terms on the RHS of Eq. 56 as the probability of acceptance of a Monte Carlo switch of the labels A and B (and thus of the controlling ‘hamiltonian’) at a point in $\{u\}$ -space, averaged over that space. In the numerator this switch is from \mathcal{E}_B to \mathcal{E}_A and the average is with respect to the canonical distribution in ensemble B ; in the denominator the roles of A and B are reversed. This is Bennett’s acceptance ratio formula [61].

Eq. 55 shows that for this strategy to work one needs the two ensembles to ‘overlap’ in the sense (somewhat less restrictive than in the case of the Zwanzig formula, Eq. 54) that the two single-ensemble PDFs are measurable at some *common* value of \mathcal{M} , the most obvious candidate being the $\mathcal{M} \simeq 0$ region intermediate between the values typical of the two ensembles. Eq. 56 shows that this requirement is effectively equivalent to the condition that the probabilities of acceptance of a hamiltonian switch can be measured, in *both* directions.

In the form described, this strategy will virtually always fail; to produce a generally workable strategy we need to introduce two further ingredients, the first essential, the second desirable. First we need to invoke ES techniques to extend the \mathcal{M} -ranges sampled in the single-ensemble simulations until they overlap; in principle this is enough to allow

us to determine the desired ratio by matching up the two distributions in Eq. 55. But the information thus gathered is more fully and efficiently utilized by taking a further step. In the $\mathcal{M} \simeq 0$ regions then accessed ‘switches’ (between the ‘hamiltonians’ of the two ensembles) can be implemented not just *virtually*, as envisaged in the acceptance ratio formula, but *actually* within a configuration space enlarged to include the macrostate label explicitly. One may then measure the full canonical PDF of $\mathcal{M} \equiv \mathcal{M}_{AB}$ and deduce the desired ratio from (Eq 50)

$$\mathcal{R}_{AB} = \frac{\int_{\mathcal{M} < 0} d\mathcal{M} P_0(\mathcal{M}|c)}{\int_{\mathcal{M} > 0} d\mathcal{M} P_0(\mathcal{M}|c)} \quad (57)$$

where we have used the fact that the sign of the order parameter acts as a signature of the ensemble.

This is the ESPS strategy. We have expressed it in a general way, as a switch between *any* nominated pair of ‘macrostates’ or ‘ensembles’. In summarizing it, we revert to the particular context of primary interest in which the two macrostates belong to *different* phases. The core idea is simple: to use ES methods to seek out regions of the configuration space of one phase which are such that a transition (switch, leap) to the other will be accepted with reasonable probability. The leap avoids mixed-phase configurations: the simulation explores both configuration spaces but is always to be found in one or the other.

The full machinery of ESPS is customizable in a number of respects, notably the choice of reference configurations, of order parameter and of transformation matrix. These issues are best explored in the context of specific examples.

2. Examples

We consider three examples of ESPS, ordered conceptually rather than chronologically.

Example A: hcp and fcc phases of Lennard Jones systems

In the case of crystalline phases it is natural to choose the reference configurations to represent the states of perfect crystalline order, described by the appropriate sets of lattice vectors

$$\{q\}_{\gamma}^{\text{REF}} \longrightarrow \{\vec{R}\}_{\gamma} \quad \gamma = \alpha, \tilde{\alpha} \quad (58)$$

Note that the particle and lattice site indexing hidden in this notation mandates a one-one mapping between the lattice sites of the two systems; we are free to choose any of the $N!$ possibilities. In the case of *hcp* and *fcc* lattices [66] it is natural to exploit the fact that the one lattice can be made out of the other merely by translating close-packed planes, in the fashion depicted in Fig. 6.

The simplest choice for the τ operations in Eq. 44 is to take them as the identity operation; the operation \mathbf{S} (Eq. 49) is then also the identity.

The choices of $\{q\}^{\text{REF}}$ and τ together define the geometry of the switch operation - the way in which a configuration of one structure is used to generate a configuration of the other: here the switch exchanges one lattice for another, while conserving the physical displacements with respect to lattice sites.

The final choice to be made is the form of the order parameter; in this case the default (built out of the energy function) defined in Eq. 53 proves the right choice. Thus, making the phase labels explicit we take

$$\mathcal{M}_{\alpha\tilde{\alpha}} = \mathcal{E}_{\alpha}(\{u\}, c) - \mathcal{E}_{\tilde{\alpha}}(\{u\}, c) \quad (59)$$

Figures 7 and 8 show results for the Lennard Jones (LJ) crystalline phases established with ESIT, on the basis of these choices [57]. Commentary (of this and other results here) is deferred to the ‘critique’ below.

Example B: hcp and fcc phases of hard spheres

We have already noted that the order parameter in ESPS need not be constructed out of the true energy functions of the two phases. The defining characteristic of an ESPS order parameter is that it measures the difference between the values of some chosen function of the common coordinate set $\{u\}$ evaluated in the two phases, such that for some region (typically ‘sufficiently small values’) of that quantity, an inter-phase switch can be successfully initiated. An ESPS ‘order parameter’ merely provides a convenient thread that can be followed to the wormhole ends.

In the case of hard spheres the energy function does not provide a usefully-graded measure of how far we are from a wormhole: the order parameter defined in Eq. 53 will generally be infinite because of hard-sphere overlap in the configurations created by the switch. But it is easy to find an alternative: all one has to do is build an order parameter out of a count of the number N^o of overlapping spheres. Instead of Eq. 59 we then have

$$\mathcal{M}_{\alpha\bar{\alpha}} \equiv N_{\alpha}^o(\{u\} c) - N_{\bar{\alpha}}^o(\{u\} c) \quad (60)$$

With the switch geometry chosen as for the LJ systems discussed above the difference between the free energies of fcc and hcp hard sphere systems can be determined precisely and transparently [56,48]. Some of the results are included in Table I.

Example C: liquid and fcc phases of hard spheres

The ESPS strategy can also be applied when one of the phases is a liquid [58]. A configuration selected at random from those explored in canonical sampling of the liquid phase will serve as a reference state. Since the liquid and solid phases generally have significantly different densities the simulation must be conducted at constant pressure [67]; the coordinate set $\{q\}$ then contains the system volume and the switch must accommodate an appropriate dilation (and can do so easily through the specification of the volumes implicit in the reference configurations). While the overlap order parameter defined in Eq. 60 remains appropriate for simulations conducted in the solid phase, in the liquid phase it is necessary to engineer something a little more elaborate to account for the fact that the particles are not spatially localized. Such considerations also lead to some relatively subtle but significant finite-size effects. Figure 9 shows some results locating the freezing pressure of hard spheres this way [58].

3. Critique

The ESPS method draws on and synthesizes a number of ideas in the extensive free energy literature, including the importance of *representations* and *space transformations between them* [63,68,69]; the utility of *expanded ensembles* in turning virtual transitions into real ones [23]; and the general power of *multicanonical methods* to seek out macrostates with any desired property [27].

Methodologically ESPS has much in common with ESIT: like ESIT it utilizes the paraphernalia of extended sampling to visit both phases in a single simulation; but in contrast to ESIT it contrives to do this *without* having to traverse the interfacial configurations which make ESIT hard, probably impossible, to implement in problems involving solid phases.

ESPS thus shares a number of the advantages that ESIT has with respect to integration methods (Sections IV A 2 IV C 2). It is pleasingly transparent: the evolution with temperature of the relative stability of *fcc* and *hcp* LJ crystals can be ‘read off’ from Fig. 7; and the LJ freezing pressure ‘seen’ in Fig. 9. Apart from finite-size effects uncertainties are purely statistical. The fact that both phases are realized within the same simulation means that finite-size effects can be handled more systematically; this seems to be a particular advantage of the ESPS approach to the liquid-solid phase boundary.

ESPS remains a computationally intensive strategy, though not prohibitively so on the scale of its competitors: one explicit comparison (in the case hard sphere crystals) indicates that ESPS and NIRM deliver similar precision for similar compute resource [34].

But the possibility of substantial improvements to ESPS remains. The idea of improved (‘targeted’) mappings between the configuration spaces has been discussed in general terms by Jarzynski [70], albeit in the context of the Zwanzig formula (Eq. 54) which will not generally work without ES-props. In the ESPS framework that mapping enters in the matrix \mathbf{S} , reflecting the representations chosen for the two phases (Eqs. 49, 44). One does not *have* to preserve the physical displacements in the course of the switch. By appropriate choice of the operations τ it is possible [71] to implement a switch which, instead, conserves a set of Fourier coordinates, and thence the *harmonic* contributions to the energy of the configurations of each phase; the determinant in Eq. 47 then captures the harmonic contribution to the free energy difference, leaving the computational problem focused on the anharmonic contributions which (alone) are left in \mathcal{R} . This strategy greatly enhances the overlap between the two branches of the order parameter distribution; but the associated efficiency gains (resulting from the reduced length of ‘path’ through \mathcal{M} -space) are offset by the greatly increased computational cost of the mapping itself [71].

V. DOING IT WITHOUT PATHS: ALTERNATIVE STRATEGIES

In this section we survey some of the strategic approaches to the phase-coexistence problem which do not fit comfortably into the path-based perspectives we have favored here.

A. The Gibbs Ensemble Monte Carlo method

1. Strategy

Gibbs Ensemble Monte Carlo (GEMC) is an ingenious method introduced by Panagiotopoulos [72], which allows one to simulate the coexistence of liquid and vapor phases without having to deal with a physical interface between them.

GEMC utilizes *two* simulation subsystems (‘boxes’); though physically separate the two boxes are thermodynamically coupled through the MC algorithm which allows them to exchange both volume and particles subject to the constraint that the total volume and number of particles remain fixed. Implementing these updates (in a way that respects detailed balance) ensures that the two systems will come to equilibrium at a common temperature, pressure and chemical potential. The temperature is fixed explicitly in the MC procedure; but the procedure itself selects the chemical potential and pressure that will secure equilibrium.

If the overall number density and temperature are chosen to lie within the two-phase region, the system must phase separate; it does so through configurations in which each box houses one pure phase, since such arrangements avoid the free energy cost of an interface. The coexistence densities can then be simply measured through a sample average over each box. By conducting simulations at a series of temperatures, the phase diagram in the temperature-density plane can be constructed. Details of the implementation procedure can be found in reference [29].

2. Critique

The GEMC method is elegant in concept and simple in practice; it seems fair to say that it revolutionized simulations of fluid phase equilibria; it has been very widely used and comprehensively reviewed [73]. We note three respects in which it is less than ideal.

First, in common with any simulation of open systems it runs into increasing difficulties as one moves down the coexistence curve to the region of high densities where particle insertion (entailed by particle exchange) has a low acceptance probability.

It also runs into difficulties of a different kind at the other end of the coexistence curve, as one approaches the critical point. GEMC effectively supposes that criticality may be identified by the coalescence of the two peaks in the separate branches of the density distribution captured by the two simulation boxes. The limiting critical behavior of the full density distribution in a system of finite size (Fig. 10, to be discussed below) shows that this is not so; the critical point cannot be reliably located this way. These difficulties are reflected in the strong finite-size-dependence of the shape of the coexistence curve evident in GEMC studies [74]. To make sense of the GEMC behavior near the critical point, therefore, one needs to invoke finite-size scaling strategies [75,76]; but these are substantially less easy to implement than they are in the framework of the μVT ensemble, to be discussed in section VI C.

Finally we note that the efficiency of GEMC is reduced through the high computational cost of volume moves (each one of which requires a recalculation of *all* inter-particle interactions). Comparisons show [73] that the strategy discussed in Sec IV C 1 (multicanonical methods and histogram re-weighting, within a grand canonical ensemble) gives a better return in terms of precision per computational unit cost. But GEMC is undoubtedly easier to implement.

B. The NPT & test particle method

1. Strategy

The NPT-TP method [77] locates phase coexistence at a prescribed temperature by finding that value of the pressure for which the chemical potentials of the two phases are equal. The chemical potential μ is identified with the difference between the Helmholtz free energies of systems containing N and $N - 1$ particles. Then the Zwanzig formula ([62], Eq. 54) shows that

$$\mu = \ln\langle e^{-[\mathcal{E}_N - \mathcal{E}_{N-1}]} \rangle_{N-1} = \mu^{ig} + \ln\langle e^{-\Delta U_{N,N-1}} \rangle_{N-1} \quad (61)$$

where $\Delta U_{N,N-1}$ is the additional configuration energy associated with the insertion of a ‘test’ particle into a system of $N - 1$ particles. Equation 61 is due to Widom [78] and forms the basis of the *test particle insertion method*. In contrast to other realizations of the Zwanzig formula it works (or can do so) reasonably well, because the argument of the exponential is of $O[1]$ rather than $O[N]$, as long as the particle interactions are short-ranged.

The values of the chemical potentials in each phase, together with their pressure-derivatives (available through the measured number densities) can be exploited to home in on the coexistence pressure. The method has been successfully applied to calculate the phase diagrams of a number of simple fluids and fluid mixtures [79,80].

2. Critique

The NPT-TP method is obviously designed to deal with the coexistence of fluid phases: particle-insertion into ordered structures is generally to be avoided. In this restricted context it is straightforward to implement, needing no more than the conventional apparatus of single-phase NPT simulation. However in common with other strategies that involve particle insertion (such as Gibbs Duhem integration, Sec. VIB) it runs into difficulties at high densities; and it cannot readily handle the behavior in the critical region (Sec. VIC 2).

C. Beyond equilibrium sampling: fast growth methods

1. Strategy

The techniques discussed in this section might reasonably have been included in our collection of path-based strategies (Section IV). However, although the idea of a ‘path’ features here too, it does so *without* the usual implications of equilibrium sampling.

At the heart of the techniques in question (we shall refer to them collectively as *Fast Growth*, FG) is a simple and beautiful result established by Jarzynski [82] which we write in the form [6]

$$\mathcal{R}_{AB} \equiv e^{-\Delta \mathcal{F}_{AB}} = \overline{e^{-W_{AB}^{\tau_s}}} \quad (62)$$

Here $W_{AB}^{\tau_s}$ is the work done in switching the effective energy function (through some time-dependent control parameter $\lambda(t)$) from \mathcal{E}_B to \mathcal{E}_A , in time τ_s , while the system observes the dynamical or stochastic updating rules appropriate to the energy function appropriate at any instant. The bar denotes an average over the ensemble of such procedures generated by choosing the initiating microstate randomly from macrostate B .

Equation 62 incorporates two more familiar claims as special cases. In the limit of *long* switching times the system has time to equilibrate at every stage of the switching procedure; then Eq. 62 reduces to the result of numerical integration along the path prescribed by the switching operation (cf Eq. 37)

$$\Delta \mathcal{F}_{AB} = \int_B^A d\lambda \langle \frac{\partial \mathcal{E}}{\partial \lambda} \rangle_\lambda \quad (63)$$

At the other extreme, if the switching time τ_s is short, the work done is just the energy cost of an instantaneous and complete hamiltonian switch, and one recovers the Zwanzig formula (cf Eq. 54)

$$\Delta \mathcal{F}_{AB} = -\ln\langle e^{-[\mathcal{E}_A - \mathcal{E}_B]} \rangle_B \quad (64)$$

In fact Eq. 62 holds *irrespective* of the switching time τ_s : the *equilibrium* free-energy difference is determined by the spectrum of a quantity, W_{AB} , associated with a *non-equilibrium* process. The ‘exponential average’ of the work done in taking the system between the designated macrostates (at *any* chosen rate) thus provides an alternative estimator of the difference between the associated free energies.

The fact that Eq. 62 (in general) folds in non-equilibrium processes is made explicit through a third result which may be deduced from it. The convexity of the exponential [81] implies that

$$\overline{e^{-W_{AB}}} \geq e^{-\overline{W_{AB}}} \quad (65)$$

so that, from Eq. 62,

$$\Delta \mathcal{F}_{AB} \leq \overline{W_{AB}} \quad (66)$$

This is the Helmholtz inequality, a variant of the Second Law and thus an acknowledgment of the consequences of irreversible processes.

2. Critique

Our discussion of the Zwanzig formula shows that the FG representation is unlikely to be practically helpful if one chooses small τ_s –at least for the systems having large enough N to be of interest to us here. But given a choice of devoting a specified computational resource to *one long* switch (and appealing to the integration procedure, Eq. 63) or *several short* switches (and appealing to the exponential averaging procedure Eq. 62) the latter seems to be the preferred strategy [82]: the two approaches are comparable in precision; but FG generates an estimate of its own uncertainties; and it is trivially implementable in parallel computing architectures. This seems a potentially fruitful avenue for further exploration.

One can avoid the issues associated with exponential-averaging if one uses the FG formula in the form of the inequality Eq 66, in tandem with the corresponding inequality emerging from a reverse switch operation (from A to B). The two results together give upper and lower bounds on the difference between the two free energies; the bounds can be tightened by a variational procedure with respect to the parameters of the chosen switch [83]. Figure 11 shows the results of such a procedure [84] applied to the Bain switch operation [85] which maps an *fcc* lattice onto a *bcc* lattice by a continuous [86] deformation.

VI. DETERMINING THE PHASE BOUNDARY: EXTRAPOLATION, TRACKING AND THE THERMODYNAMIC LIMIT

The path-based methods exemplified in the preceding section provide us with ways of estimating the difference between the free energies of the two phases, $\Delta\mathcal{F}_{\alpha\tilde{\alpha}}(\{\lambda\}, N)$ (Eq. 10) at *some* point $c \equiv \{\lambda\}$ in the space of the controlling fields [88], for a system whose size N (which we shall occasionally make explicit in this section) is computationally manageable.

The practical task of interest here requires that we identify the set of points $c_x \equiv \{\lambda\}_x$ giving phase coexistence in the thermodynamic limit, and defined by the solutions to the equation

$$\lim_{N \rightarrow \infty} \frac{1}{N} \Delta\mathcal{F}_{\alpha\tilde{\alpha}}(\{\lambda\}_x, N) = 0 \quad (67)$$

We divide this programme into three parts. The first issue is how to use the data accumulated at our chosen state point to infer the location of *some* point c_x on the coexistence curve, for some finite N . The second is to map out the phase boundary emanating from that point. And the third is to deal with the corrections associated with the limited (‘finite’) size of the simulation system.

We shall restrict the discussion to a two-dimensional field space spanned by fields $\{\lambda\} \equiv \lambda_1, \lambda_2$ with conjugate macrovariables $\{M\} \equiv M_1, M_2$.

A. Extrapolation to the phase boundary

The simplest way of using the measurements at $\{\lambda\}$ to estimate the location of a point on the phase boundary is to perform a linear extrapolation in one of the fields using the result

$$\frac{\partial \Delta\mathcal{F}_{\alpha\tilde{\alpha}}(\{\lambda\})}{\partial \lambda} = \langle M \rangle_\alpha - \langle M \rangle_{\tilde{\alpha}} \quad (68)$$

Note that the quantities on the RHS of this equation represent separate single-phase expectation values, defined with respect to single-phase canonical distributions of the form given in Eq. 4; they are thus problem-free.

The implied estimate of the coexistence-value of the field λ_1 (say) is then

$$\lambda_{1x} = \lambda_1 - \frac{\Delta\mathcal{F}_{\alpha\tilde{\alpha}}(\{\lambda\})}{\langle M_1 \rangle_\alpha - \langle M_1 \rangle_{\tilde{\alpha}}} \quad (69)$$

Such extrapolations provide a simple, but possibly crude, way of correcting an initially poor estimate of coexistence.

One may be able to do better by appealing to histogram re-weighting techniques (Appendix B). *If* the initial measurement of a free energy difference is based on some form of extended sampling which establishes the full canonical distribution of some order parameter; and *if* that order parameter is the conjugate of one of the fields spanning the phase diagram of interest; *then* HR allows one to scan through a range of values of that field to find the coexistence value (identified by the resulting equality of the areas of the two peaks in the canonical order parameter distribution, implied by Eq. 10). Such a process is easily automated.

B. Tracking the phase boundary

In principle knowledge of a *single point* on the coexistence curve permits the *entire curve* to be traced without further calculation of free energies. The key result needed follows simply from Eq. 68, applied to each field λ_1 and λ_2 in turn. The *slope* of the coexistence curve $\Delta\mathcal{F}_{\alpha\bar{\alpha}}(\{\lambda\}_x) = 0$, follows as

$$\left[\frac{d\lambda_1}{d\lambda_2}\right]_x = -\frac{\langle M_2 \rangle_\alpha - \langle M_2 \rangle_{\bar{\alpha}}}{\langle M_1 \rangle_\alpha - \langle M_1 \rangle_{\bar{\alpha}}} \quad (70)$$

This is the generalized Clausius-Clapeyron equation [89]. It expresses the slope entirely in terms of single-phase averages; the slope can be employed (in a predictor-corrector scheme) to estimate a nearby coexistence point. Fresh simulations performed at this new point yield the phase boundary gradient there, allowing further extrapolation to be made, and so on. In this manner one can in principle track the whole coexistence curve. This strategy is widely known as Gibbs-Duhem integration (GDI) [90].

GDI has been used effectively in a number of studies, most notably in the context of freezing of hard and soft spheres [91]. Its distinctive feature is simultaneously its strength and its weakness: once boot-strapped by knowledge of one point on the coexistence curve it subsequently requires only *single-phase* averages. This is clearly a virtue since the elaborate machinery needed for two-phase sampling is not to be unleashed lightly. But *without* any ‘reconnection’ of the two configuration-spaces at subsequent simulation state points, the GDI approach offers no feedback on integration errors. Since there will generally exist a *band* of practically stable states on each side of the phase boundary, it is possible for the integration to wander significantly from the true boundary with no indication that anything is wrong.

A more robust (though computationally more intensive) alternative to GDI is provided by a synthesis of extended (multicanonical) sampling and histogram re-weighting techniques. The method is boot-strapped by an ES measurement of the full canonical distribution of a suitable order parameter, at some point on the coexistence curve (identified by the equal areas criterion specified in Eq. 10). HR techniques then allow one to map a region of the phase boundary close to this point. The range over which such extrapolations are reliable is limited (Appendix B) and it is not possible to extrapolate arbitrarily far along the phase boundary: further multicanonical simulations will be needed at points that lie at the extremes of the range of reliable extrapolation. But there is no need to determine a new set of weights (a new extended sampling distribution) from *scratch* for these new simulations. HR allows one to generate a *rough estimate* of the equilibrium order parameter (at these points) by extrapolation from the original measured distribution. The ‘rough estimate’ is enough to furnish a usable set of weights for the new multicanonical simulations. Repeating the combined procedure (multicanonical simulation followed by histogram extrapolation) one can track along the coexistence curve. The data from the separate histograms can subsequently be combined self consistently (through multihistogram extrapolation, as discussed in Appendix B) to yield the whole phase boundary. If one wishes to implement this procedure for a phase boundary that terminates in a critical point it is advisable to start the tracking procedure nearby. At such a point the ergodic block presented to inter-phase traverses is relatively small (the canonical order parameter distribution is relatively weakly doubly-peaked); and so the multicanonical distribution (weights) required to initiate the whole process can be determined without extensive (perhaps without any) iterative procedures [92].

C. Finite-size effects

Computer simulation is invariably conducted on a model system whose size is small on the thermodynamic scale one typically has in mind when one refers to ‘phase diagrams’. Any simulation-based study of phase behavior thus necessarily requires careful consideration of ‘finite-size effects’. The nature of these effects is significantly different according to whether one is concerned with behavior close to or remote from a critical point. The distinction reflects the relative sizes of the linear dimension L of the system –the edge of the simulation cube– and the correlation length ξ –the distance over which the local configurational variables are correlated. By ‘non-critical’ we mean a system for which $L \gg \xi$; by critical we mean one for which $L \ll \xi$. We shall discuss these two regions in turn, and avoid the lacuna in between.

1. Non-critical systems

In the case of non-critical systems the issue of finite-size effects is traditionally expressed in terms of the finite-size corrections to the free energy densities of each of the two phases:

$$f_\alpha(c, \infty) - f_\alpha(c, N) \quad \text{where} \quad f_\alpha(c, N) = \frac{1}{N} \mathcal{F}_\alpha(c, N) \quad (71)$$

In recent years substantial efforts have been made to develop a theoretical framework for understanding the nature of such corrections [93]. In the case of *lattice models* (ie models of strictly localized particles) in the *NVT-ensemble* with *periodic boundary conditions* (PBCs) it has been established *a priori* [94] and corroborated in explicit simulation [95] that the corrections are *exponentially small* in the system size [96].

However these results do not immediately carry over to the problems of interest here where (while PBCs are the norm) the ensembles are frequently open or constant pressure, and the systems do not fit in to the lattice model framework. Even in the apparently simple case of crystalline solids in *NVT*, the free translation of the center of mass introduces N -dependent phase space-factors in the configurational integral which manifest themselves as additional finite-size corrections to the free energy; these may not yet be fully understood [97,58]. If one adopts the traditional stance then, one is typically faced with having to make extrapolations of the free energy densities in *each* of the two phases, *without* a secure understanding of the underlying form ($\frac{1}{N}$? $\frac{\ln N}{N}$? ...) of the corrections involved.

The problems are reduced if one shifts the focus of attention from the single-phase free energies to the quantities of real interest: the difference between the two free energies, and the field-values that identify where it vanishes. In both cases within an ES strategy that treats both phases together, there is only *one* extrapolation to do, which is clearly a step forward. If the two phases are of the same generic type (eg two crystalline solids) one can expect cancellations of corrections of the form $\frac{\ln N}{N}$ –which should be identical in both phases– leaving presumably at most $\frac{1}{N}$ corrections. If the phases are not of the same generic type (a solid and a liquid) the logarithmic corrections will probably not cancel; reliable extrapolation will be possible only if they can be identified and allowed for explicitly, leaving only a pure power law.

Overall it seems that there is considerable room for progress here.

2. Near-critical systems

In the case of simulation studies of near-critical systems, the issues associated with ‘finite-size effects’ have an altogether different flavor. First, they are no longer properly regarded as essentially small effects to be ‘corrected for’; the critical region is characterized by a strong and distinctive dependence of system-properties on system size; the right strategy is to address that dependence head on, and exploit it. Second, there is an extensive framework to appeal to here: the phenomenology of *finite-size scaling* [93] and the underpinning theoretical structure of the *renormalization group* [98] together show what to look for and what to expect. Third, the objectives are rather different, going well beyond the issues of phase-diagram mapping that we are preoccupied with here.

Our discussion will be substantially briefer than it might be; we will focus on the issue most relevant here (but not not the most interesting) – the location of the critical point in a fluid.

As in the coexistence curve problem, the key is the distribution of the order parameter, in this case the density. On the coexistence curve, remote from the critical point (ie in the region $\xi \ll L$) we have seen (Fig. 5) that this distribution comprises two peaks of equal area, each roughly centered on the corresponding single phase average; the two peaks are narrow and near-Gaussian in form (the more so the larger the system size) as one would expect from the Central Limit Theorem; the probability of inter-phase tunneling (an inverse measure of the ergodic barrier) is vanishingly small [50]. As one moves up the coexistence curve simulations show more or less what one would guess: the peaks broaden; they become less convincingly Gaussian; and the tunneling probability increases: this is the natural evolution en route to the form which must be appropriate in the one-phase region beyond criticality – a single peak narrowing with increasing L , asymptotically Gaussian. Against this immediately intelligible backdrop, the behavior at criticality (Fig. 10) is comparatively subtle. Again (for large enough L , but still $L \ll \xi$) a limiting form is reached; however that limiting form comprises neither one Gaussian nor two, but something in between. The distribution narrows with increasing L (while preserving the shape of the ‘limiting form’); however here the width varies not as $1/L^{d/2} = 1/\sqrt{N}$ (familiar from Central Limit behavior) but as $1/L^{\beta/\nu}$ where β and ν are the critical exponents of the order parameter and the correlation length respectively. There is good reason to believe that the shape of the distribution shares the distinctive quality of the critical exponents – all are ‘universal signatures of behavior common to a wide range of physically disparate systems. The idea (of long-standing [99]) that fluids and Ising magnets belong to the same universality class is corroborated (beyond the exponent level) by the correspondence between their critical-point order parameter distributions [100]. Once that correspondence has been established to ones satisfaction one is at liberty to *exploit* it to refine the assignment of fluid critical-point parameters. The form of the density distribution depends sensitively on the choice for the controlling fields (chemical potential μ and temperature T); *demanding* correspondence with the form appropriate for the Ising universality class (studied extensively and known with considerable precision) sets narrow bounds on the location of the fluid critical point [44].

Finally we should not overlook one -perhaps unexpected- feature of the critical distribution: it shows that clear signatures of the two (incipient) phases persist along the coexistence curve and *right through to the critical point* in a finite system. Failure to appreciate this will lead (as it has in the past) to substantial overestimates of critical temperatures.

VII. DEALING WITH IMPERFECTION

Real substances often deviate from the idealized models employed in simulation studies. For instance many complex fluids, whether natural or synthetic in origin, comprise mixtures of *similar* rather than *identical* constituents. Similarly, crystalline phases usually exhibit a finite concentration of defects which disturb the otherwise perfect crystalline order. The presence of imperfections can significantly alter phase behavior with respect to the idealized case. If one is to realize the goal of obtaining quantitatively accurate simulation data for real substances, the effects of imperfections must be incorporated. In this section we consider the state-of-the-art in dealing with two kinds of imperfections, polydispersity and point defects in crystals.

A. Polydispersity

Statistical mechanics was originally formulated to describe the properties of systems of *identical* particles such as atoms or small molecules. However, many materials of industrial and commercial importance do not fit neatly into this framework. For example, the particles in a colloidal suspension are never strictly identical to one another, but have a range of radii (and possibly surface charges, shapes etc). This dependence of the particle properties on one or more continuous parameters is known as polydispersity. One can regard a polydisperse fluid as a mixture of an infinite number of distinct particle species. If we label each species according to the value of its polydisperse attribute, σ , the state of a polydisperse system entails specification of a density *distribution* $\rho(\sigma)$, rather than a finite number of density variables. It is usual to identify two distinct types of polydispersity: *variable* and *fixed*. Variable polydispersity pertains to systems such as ionic micelles or oil-water emulsions, where the degree of polydispersity (as measured by the form of $\rho(\sigma)$) can change under the influence of external factors. A more common situation is fixed polydispersity, appropriate for the description of systems such as colloidal dispersions, liquid crystals and polymers. Here the form of $\rho(\sigma)$ is determined by the synthesis of the fluid.

Computationally, polydispersity is best handled within a grand canonical (GCE) or semi-grand canonical ensemble in which the density distribution $\rho(\sigma)$ is controlled by a conjugate chemical potential distribution $\mu(\sigma)$. Use of such an ensemble is attractive because it allows $\rho(\sigma)$ to fluctuate as a whole, thereby sampling many different realizations of the disorder and hence reducing finite-size effects. Within such a framework, the case of variable polydispersity is considerably easier to tackle than fixed polydispersity: the phase behavior is simply obtained as a function of the width of the prescribed $\mu(\sigma)$ distribution. Perhaps for this reason, most simulation studies of phase behavior in polydisperse systems have focused on the variable case [101,90,102,103].

Handling fixed polydispersity is computationally much more challenging: one wishes to retain the efficiency of the GCE, but to do so, a way must be found to adapt the imposed form of $\mu(\sigma)$ such as to realize the prescribed form of $\rho(\sigma)$. This task is complicated by the fact that $\rho(\sigma)$ is a *functional* of $\mu(\sigma)$. Recently, however, a new approach has been developed which handles this difficulty. The key idea is that the required form of $\mu(\sigma)$ is obtainable iteratively by functionally minimizing a cost function quantifying the deviation of the measured form of $\rho(\sigma)$ from the prescribed ‘target’ form. For efficiency reasons, this minimization is embedded within a histogram re-weighting scheme (Appendix B) obviating the need for a new simulation at each iteration. The new method is efficient, as evidenced by tests on polydisperse hard spheres [104] where it permitted the first direct simulation measurements of the equation of state of a polydisperse fluid.

B. Crystalline defects

Defects in crystals are known to have a potentially major influence on phase behavior. For instance, dislocation unbinding is believed to be central to the 2D melting transition, while in 3D there is evidence to suggest that defects can act as nucleation centers for the liquid phase [105]. In superconductors, defects can pin vortices and influence ‘vortex melting’ [106]

Almost all computational studies of defect free-energies (and their influence on phase transitions) have been concerned with point defects. Polson *et al* [97], used the results of early calculations [107] of the vacancy free energy of a

hard sphere crystal, to estimate the equilibrium vacancy concentration at melting. Comparison with the measured free energies of the perfect hard sphere crystal (obtained from the Einstein Crystal NIRM method discussed in Sec IV A) allowed them to estimate the effect of vacancies on the melting pressure, predicting a significant shift. A separate calculation for interstitials found their equilibrium concentration at melting to be 3 orders of magnitude smaller than that of vacancies. In follow-up work, Pronk and Frenkel [108] used a technique similar to the Widom particle insertion method (Section V B 1) to calculate the vacancy free energy of a hard sphere crystal. For interstitial defects they employed an extended sampling technique in which a tagged ‘ghost particle’ is grown reversibly in an interstitial site.

To our knowledge there have been no reported measurements of equilibrium defect concentrations in soft spheres models. Similarly, relatively few measurements have been reported of defect free energies in models for real systems. Those that exist rely on integration methods to connect the defective solid to the perfect solid. In *ab-initio* studies the computational cost of this procedure can be high, although results have recently started to appear, most notably for vacancies and interstitial defects in Silicon. For a review see reference [109].

VIII. OUTLOOK

Those who read this paper may share with its authors the feelings expressed in reference [110]: the dynamics in this particular problem space seems to have been rather more diffusive than ballistic. It is therefore wise to have some idea of where the ultimate destination is, and the strategies that are most likely to take us there.

The long term goal is a computational framework that will be grounded in electronic structure as distinct from phenomenological particle potentials; that will predict global phase behavior *a priori*, rather than simply decide between two nominated candidate phases; and that will handle quantum behavior, in contrast to the essentially classical framework on which we have focused here. That goal is distant but not altogether out of sight. Integrating *ab-initio* electronic structure calculations with the statistical mechanics of phase behavior has already received some attention [109,111]. The WL algorithm ([112], Appendix A) offers a glimpse of the kind of self-monitoring configuration-space search algorithm that one needs to make automated *a priori* predictions of phase behavior possible. And folding in quantum mechanics requires only a dimensionality upgrade [113].

There are of course many other challenges, a little less grand: the two space and time scales arising in asymmetrical binary mixtures [114]; the fast attrition (exponential in the chain length) for the insertion of polymers in NPT-TP or grand-canonical methods [115]; the long range interactions in coulombic fluids [116]; the extended equilibration times for dense liquids near the structural glass transition [117]; and the extreme long-time-dynamics of the escape from metastable states in nanoscale ferromagnets [118].

As regards the strategies that seem most likely to take us forward, we make three general observations:

Making the most of the information available in simulation studies requires an understanding of finite-size effects; it also requires awareness of the utility of quantities that one would not naturally consider were one restricted to pen and paper.

We will surely need new algorithms; they come from physical insight into the configurational core of the problem at hand.

One needs to match formulations to the available technology: parallel computing architectures give some algorithms a head-start.

APPENDIX A: BUILDING EXTENDED SAMPLING DISTRIBUTIONS

In contrast to canonical sampling distributions whose form can be *written down* (Eq. 1) the Extended Sampling (ES) distributions discussed in section III C have to be *built*. There is a large literature devoted to the building techniques, extending back at least as far as reference [22]. We restrict our attention to relatively recent developments (those that seem to be reflected in current practices); and we shall focus on those aspects which are most relevant to ES distributions facilitating two-phase sampling.

In the broad-brush classification scheme offered in Sec. III C the domain of an ES distribution may be prescribed by a range of values of one or more *fields*, or one or more *macrovariables*. We shall focus on the latter representation which seems simpler to manage. The generic task then is to construct a (‘multicanonical’) sampling distribution which will visit all (equal-sized) intervals within a chosen range of the nominated macrovariable(s) with roughly equal probability: the multicanonical distribution of the macrovariable(s) is essentially flat over the chosen range.

In formulating a strategy for building such a distribution, most authors have chosen to consider the particular case in which the macrovariable space is one-dimensional, and is spanned by the configurational *energy*, E [119]. The choice is motivated by the fact that a distribution that is multicanonical in E samples configurations typical of a range of

temperatures, providing access to the simple (reference-state) behavior that often sets in at high or low temperatures. For the purposes of two-phase sampling we typically need to track a path defined on some macrovariable other than the energy (ideally, in *addition* to it: we will come back to this). The hallmarks of a ‘good’ choice are that in some region of the chosen variable (inevitably one with intrinsically low equilibrium probability) the system may pass (has a workably-large chance of passing) from one phase to the other. In discussing the key issues, then, we shall have in mind this kind of quantity; we shall continue to refer to it as an order parameter, and denote it by M .

One can easily identify the generic structure of the sampling distribution we require. It must be of the form

$$P_S(\{q\}) \doteq \frac{P_0(\{q\} | c)}{\hat{P}_0(M(\{q\}))} \quad (\text{A1})$$

Here $\hat{P}_0(M)$ is an *estimate* of the true canonical M -distribution. Appealing to the sampling identity 24 the resulting M -distribution is

$$P_S(M) = \langle \delta[M - M(\{q\})] \rangle_S \doteq \langle \frac{P_S}{P_0} \delta[M - M(\{q\})] \rangle_0 \doteq \frac{P_0(M | c)}{\hat{P}_0(M)} \quad (\text{A2})$$

and is multicanonical (flat) to the extent that our estimate of the canonical M -distribution is a good one.

We can also immediately write down the prescription for generating the ensemble of configurations defined by the chosen sampling distribution. We need a simple MC procedure with acceptance probability (Eq. 13, with the presumption of Eq. 16)

$$P_A(\{q\} \rightarrow \{q'\}) = \min \left\{ 1, \frac{P_S(\{q'\})}{P_S(\{q\})} \right\} \quad (\text{A3})$$

In turning this skeleton framework into a working technique one must make choices in regard to three key issues:

1. How to *parameterize* the estimator \hat{P}_0 .
2. What *statistics* of the ensemble to use to guide the update of \hat{P}_0 .
3. What *algorithm* to use in updating \hat{P}_0 .

The second and third issues are the ones of real substance; the issue of parameterization is important only because the proliferation of different choices that have been made here may give the impression that there are more techniques available than is actually the case. That proliferation is due, in some measure, to the preoccupation with building ES distributions for the *energy*, E . There are as many ‘natural parameterizations’ here as there are ways in which E appears in canonical sampling. Thus Berg and Neuhaus [27] employ an E -dependent *effective temperature*; Lee [120] utilizes a *microcanonical entropy* function; Wang and Landau [112] focus on a *density of states* function. Given our concern with macrovariables other than E the most appropriate parameterisation of the sampling distribution here is through a *multicanonical weight function* $\hat{\eta}(M)$, in practice represented by a discrete set of multicanonical weights $\{\hat{\eta}\}$. Thus we write [123]

$$\hat{P}_0(M | c) \doteq e^{-\hat{\eta}(M)} \quad (\text{A4})$$

implying (through Eq. A1) a sampling distribution

$$P_S(\{q\}) \doteq \sum_{j=1}^{\Omega} e^{-\mathcal{E}(\{q\}) - \hat{\eta}_j} \Delta_j[M] \doteq e^{-\beta E(\{q\}) + \hat{\eta}[M(\{q\})]} \quad (\text{A5})$$

which is of the general form of Eq. 32, with $w_j \equiv e^{\hat{\eta}_j}$.

There are broadly two strategic responses to the second of the issues raised above: to drive the estimator in the right direction one may appeal to the *statistics of visits to macrostates* or to the *statistics of transitions between macrostates*. We divide our discussion accordingly.

1. Statistics of visits to macrostates

The extent to which any chosen sampling distribution (weight function) meets our requirements is reflected most directly in the M -distribution it implies. One can estimate that distribution from a histogram $H(M)$ of the macrostates visited in the course of a set of MC observations. One can then use this information to refine the sampling distribution to be used in the next set of MC observations. The simplest update algorithm is of the form [122]

$$\hat{\eta}(M) \longrightarrow \hat{\eta}(M) - \ln[H(M) + 1] + k \quad (\text{A6})$$

The form of the logarithmic term serves to ensure that macrostates registering no counts (of which there will usually be many) have their weights incremented by the same finite amount (the positive constant k [124]); macrostates which *have* been visited are (comparatively) down-weighted.

Each successive iteration comprises a fresh simulation, performed using the weight function yielded by its predecessor; since the weights attached to unvisited macrostates is enhanced (by k) at every iteration which fails to reach them, the algorithm plumbs a depth of probability that grows exponentially with the iteration number. The iterations proceed until the sampling distribution is roughly flat over the entire range of interest.

There are many tricks of the trade here. One must recognize the interplay between signal and noise in the histograms: the algorithm will converge only as long as the signal is clear. To promote faster convergence one can perform a linear extrapolation from the sampled into the unsampled region. One may bootstrap the process by choosing an initial setting for the weight function on the basis of results established on a smaller (computationally-less-demanding) system. To avoid spending excessive time sampling regions in which the weight function has already been reliably determined, one can adopt a multistage approach. Here one determines the weight function separately within slightly overlapping windows of the macrovariable. The individual parts of the weight function are then synthesized using multi-histogram re-weighting (Appendix B) to obtain the full weight function. For further details the reader is referred to [121].

The strategy we have discussed is generally attributed to Berg and Neuhaus (BN) [27]. Wang and Landau (WL) have offered an alternative formulation [112]. To expose what is different, and what is not, it is helpful to consider first the case in which the macrovariable is the energy, E . Appealing to what one knows *a priori* about the canonical energy distribution, the obvious parameterization is

$$\hat{P}_0(E) \doteq \hat{G}(E)e^{-\beta E} \quad (\text{A7})$$

where $\hat{G}(E)$ is an estimator of the *density of states* function $G(E)$. Matching this parameterization to the multi-canonical weight function $\hat{\eta}(E)$ implied by choosing $M = E$ in Eq. A4 one obtains the correspondence

$$\hat{\eta}(E) = \beta E - \ln \hat{G}(E) \quad (\text{A8})$$

There is thus no major difference here. The differences between the two strategies reside rather in the procedure by which the parameters of the sampling distribution are *updated*, and the point at which that procedure is *terminated*.

Like BN, WL monitors visits to macrostates. But, while BN updates the weights of *all* macrostates after *many* MC steps, WL updates its ‘density of states’ for *the current* macrostate after *every* step. The update prescription is

$$\hat{G}(E) \longrightarrow f \hat{G}(E) \quad (\text{A9})$$

where f is a constant, greater than unity. As in BN a visit to a given macrostate tends to reduce the probability of further visits. But in WL this change takes place *immediately* so the sampling distribution evolves on the basic timescale of the simulation. As the simulation proceeds, the evolution in the sampling distribution irons out large differences in the sampling probability across E -space, which is monitored through a histogram $H(E)$. When that histogram satisfies a nominated ‘flatness criterion’ the entire process is repeated (starting from the current $\hat{G}(E)$, but zeroing $H(E)$) with a smaller value of the weight-modification factor, f .

Like BN, then, the WL strategy entails a two-time scale iterative process. But in BN the aim is only to generate a set of weights that can be utilized in a further, final multicanonical sampling process; the iterative procedure is terminated when the weights are sufficiently good to allow this. In contrast, in WL the iterative procedure is pursued further –to a point [125] where $\hat{G}(E)$ may be regarded as a definitive approximation to $G(E)$, which can be used to compute any (single phase) thermal property at any temperature through the partition function

$$Z(\beta) = \int dE G(E) e^{-\beta E} \quad (\text{A10})$$

In the context of energy sampling, then, BN and WL achieve essentially the same ends, by algorithmically-different routes. Both entail choices (in regard to their update schedule) which have to rest on experience rather than any deep understanding. WL seems closer to the self-monitoring ideal, and may scale more favorably with system size.

The WL procedure *can* be applied to any chosen macrovariable, M . But while a good estimate $\hat{G}(E)$ is sufficient to allow multicanonical sampling in E (and a definitive one is enough to determine $Z(\beta)$, Eq. A10) the M -density of states does not itself deliver the desired analogues: we need, rather, the *joint* density of states $G(E, M)$ which determines the restricted, single-phase partition functions through

$$Z_\alpha(\beta) = \int dE \int dM G(E, M) e^{-\beta E} \Delta_\alpha[M] \quad (\text{A11})$$

The WL strategy *does* readily generalize to a 2D macrovariable space. The substantially greater investment of computational resources is offset by the fact that the relative weights of the two phases can be determined at *any* temperature. Reference [126] provides one of (as yet few) illustrations of this strategy, which seems simple, powerful and general.

2. Statistics of transitions between macrostates

The principal general feature of the algorithms based on visited macrostates is that the domain of the macrovariable they explore expands relatively slowly into the regions of interest. The algorithms we now discuss offer significant improvement in this respect. Although (inevitably, it seems) they exist in a variety of guises, they have a common core which is easily established. We take the general detailed balance condition Eq. 11 and sum over configurations $\{q\}$ and $\{q'\}$ that contribute (respectively) to the macrostates M_i and M_j of some chosen macrovariable M . We obtain immediately

$$P_S(M_i) \overline{P_S(M_i \rightarrow M_j)} = P_S(M_j) \overline{P_S(M_j \rightarrow M_i)} \quad (\text{A12})$$

The terms with over-bars are macrostate transition probabilities (TP). Specifically $\overline{P_S(M_i \rightarrow M_j)}$ is the probability (per unit time, say) of a transition from some nominated configuration in M_i to *any* configuration in M_j , ensemble-averaged over the configuration in M_i . Adopting a more concise (and suggestive) matrix notation:

$$p_S^M[i] \rho_S^M[ij] = p_S^M[j] \rho_S^M[ji] \quad (\text{A13})$$

This is a not-so-detailed balance condition; it holds for any sampling distribution and any macrovariable [122]. The components of the eigenvector of the TP matrix (of eigenvalue unity) thus identify the macrostate probabilities. This is more useful than it might seem. One can build up an approximation of the transition matrix by monitoring the transitions which follow when a simulation is launched from an *arbitrary* point in configuration space. The ‘arbitrary’ point can be judiciously sited in the heart of the interesting region; the subsequent simulations then carry the macrovariable right through the chosen region, allowing one to accumulate information about it from the outset. With a sampling distribution parameterized as in Eq. A5, the update scheme is simply

$$\hat{\eta}(M) \longrightarrow \hat{\eta}(M) - \ln[\hat{P}_S(M)] + k \quad (\text{A14})$$

where $\hat{P}_S(M)$ is the estimate of the sampling distribution that is *deduced* from the measured TP matrix [127]. Reference [55] describes the application of this technique to a structural phase transition.

One particular case of Equation A12 has attracted considerable attention. If one sets $M = E$, and considers the infinite temperature limit, the probabilities of the macrostates E_i and E_j can be replaced by the associated values of the density of states function $G(E_i)$ and $G(E_j)$. The resulting equation has been christened *the broad-histogram relation* [128]; it forms the core of extensive studies of transition probability methods referred to variously as ‘flat histogram’ [129] and ‘transition matrix’ [130]. Applications of these formulations seem to have been restricted to the situation where the energy is the macrovariable, and the energy spectrum is discrete.

Methods utilizing macrostate transitions do have one notable advantage with respect to those that rely on histograms of macrostate-visits. In transition-methods the results of separate simulation runs (possibly initiated from different points in macrovariable space) can be straightforwardly combined: one simply aggregates the contributions to the transition-count matrix [122]. Synthesizing the information in separate histograms (see the succeeding Appendix) is less straightforward. The easy synthesis of data sets makes the TP method ideally suited for implementation in parallel architectures. Whether these advantages are sufficient to offset the the fact that TP methods are undoubtedly more complicated to implement is, perhaps, a matter of individual taste.

APPENDIX B: HISTOGRAM RE-WEIGHTING

Without loss of generality the effective configurational energy may always be written in the form

$$\mathcal{E}(\{q\}, \{\lambda\}) = - \sum_{\mu} \lambda^{(\mu)} M^{(\mu)}(\{q\}) \quad (\text{B1})$$

where $\{\lambda\}$ and $\{M\}$ are sets [88] comprising one or more mutually conjugate fields and macrovariables [132]. We consider two ensembles which differ only in the values of one or more of the fields $\{\lambda\}$. The canonical sampling distributions of the two ensembles are then related by

$$P_0(\{q\}|\{\lambda'\})e^{-\sum_{\mu} \lambda'^{(\mu)} M^{(\mu)}(\{q\})} \doteq P_0(\{q\}|\{\lambda\})e^{-\sum_{\mu} \lambda^{(\mu)} M^{(\mu)}(\{q\})} \quad (\text{B2})$$

where, again, \doteq signifies equality to within a configuration-independent constant. Performing the configurational sum, for *fixed* values of the macrovariables $\{M\}$ then yields the relationship

$$P_0(\{M\}|\{\lambda'\})e^{-\sum_{\mu} \lambda'^{(\mu)} M^{(\mu)}} \doteq P_0(\{M\}|\{\lambda\})e^{-\sum_{\mu} \lambda^{(\mu)} M^{(\mu)}} \quad (\text{B3})$$

so that

$$P_0(\{M\}|\{\lambda'\}) \doteq P_0(\{M\}|\{\lambda\})e^{-\sum_{\mu} [\lambda^{(\mu)} - \lambda'^{(\mu)}] M^{(\mu)}} \quad (\text{B4})$$

In principle then the $\{M\}$ -distribution for *any* values of the fields can be inferred from the $\{M\}$ -distribution for *one* particular set of values [131]. This is the basis of the Histogram Re-weighting (HR) technique [133], also known as Histogram Extrapolation [134]. It can be seen as a special case of the sampling identity given in Eq. 24 [135]. Like that identity its formal promise is not matched by what it can deliver in practice. The measurements in the $\{\lambda\}$ -ensemble (from which the extrapolation is to be made) determine only an *estimate* of the $\{M\}$ distribution for that ensemble. The estimate will be relatively good for the most probable $\{M\}$ values (around the ‘peak’ of that distribution) which are ‘well-sampled’, and relatively poor for the less probable values (in the ‘wings’), which are sampled less well. This trade-off (desirable if one wants only properties of the $\{\lambda\}$ ensemble) limits the range of field-space over which extrapolation will be reliable. The distribution associated with a set $\{\lambda'\}$ remote from λ may peak in a region of $\{M\}$ -space far from the peak of the measured distribution, lying instead in its poorly sampled wings. In such circumstances the estimate provided by the extrapolation prescribed by Eq. B4 will be unreliable (and will typically reveal itself as such, through its ragged appearance).

In the scheme we have discussed extrapolations are made on the basis of a histogram determined at a *single* point in field-space. The multi-histogram method [133] extends this framework. It entails a sequence of separate simulations spanning a range of the field (or fields) whose conjugate macrovariable(s) are of interest. The intervals are chosen so that the tails of the histograms of the macrovariable accumulated at neighboring state points overlap. It is clear from the discussion above that in principle *every* histogram will provide some information about *every* region; and that the most reliable information about any *given* region will come from the histogram which samples that region most effectively. These ideas can be expressed in an explicit prescription for synthesizing all the histograms to give an estimate of the canonical distribution of the macrovariable across the whole range of fields [133,136].

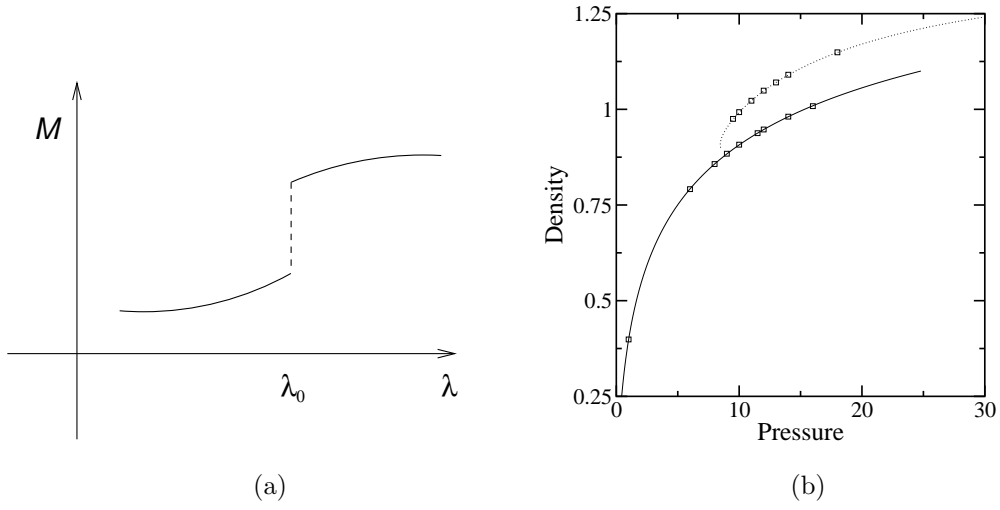


FIG. 1. (a) Schematic representation of the results of an ‘ideal experiment’ on phase behavior, which may take as much time as we need: the equilibrium value of some physical quantity M changes discontinuously at some sharply-defined value, λ_0 , of some field λ . (b) An example of the experimental reality: the results of a typical simulation study of solid-liquid phase behavior of hard spheres; the measured density continues to follow the branch (liquid or solid) on which the simulation is initiated, well beyond the coexistence pressure P_x ($\simeq 11.3$ in these units) [5].

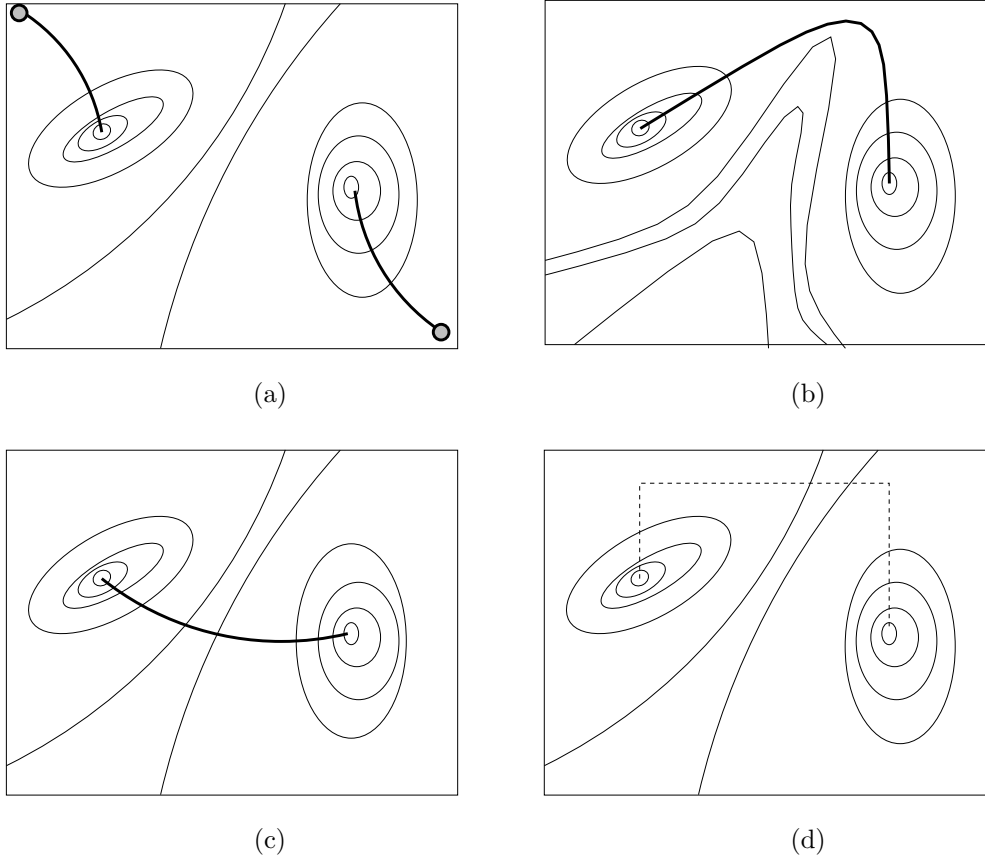


FIG. 2. Schematic representation of the four conceptually-different paths (the heavy lines) one may utilize to attack the phase-coexistence problem. Each figure depicts a configuration space spanned by two macroscopic properties (such as energy, density ...); the contours link macrostates of equal probability, for some given conditions c (such as temperature, pressure ...). The two mountain-tops locate the equilibrium macrostates associated with the two competing phases, under these conditions. They are separated by a probability-ravine (free-energy-barrier). In case (a) the path comprises two disjoint sections confined to each of the two phases and terminating in appropriate reference macrostates. In (b) the path skirts the ravine. In (c) it passes through the ravine. In (d) it leaps the ravine.




FIG. 3. Probability distribution of the number N of particles in a LJ fluid. The simulations use the HPT method described in section IV B 1. The solid line shows the distribution for a replica whose $\mu - T$ parameters lie close to coexistence; the dashed line (offset) shows the distribution (for the same $\mu - T$ parameters) obtained by folding in (explicitly) the contributions of all replicas, using multi-histogram re-weighting. (Taken from Figure 2 of Reference [38]).




FIG. 4. Phase diagram for the LJ fluid. The squares show results obtained by the HPT method described in section IV B 1 and illustrated in Figure 3. The triangles show results obtained by the ESIT strategy described in section IV C 1 and illustrated in figure 5. (Taken from Figure 3 of Reference [38]).




FIG. 5. Results from a multicanonical simulation of the 3D Lennard-Jones fluid at a point on the coexistence curve. The figure shows both the multicanonical sampling distribution $P_S(\rho)$ (symbols: \circ) and the corresponding estimate of the equilibrium distribution $P_0(\rho)$ with $\rho = N/V$ the number density. The inset shows the value of the equilibrium distribution in the interfacial region. [43].




FIG. 6. The simple transformation for switching between *fcc* and *hcp* lattices. The diagram shows 6 close-packed ($x - y$) layers. (The additional bracketed layer at the bottom is the periodic image of the layer at the top.) The circles show the boundaries of particles located at the sites of the two close-packed structures. In the lattice switch operation the top pair of planes are left unaltered, while the other pairs of planes are relocated by translations, specified by the black and white arrows. The switch operation is discrete: the relocations occur ‘through the wormholes’. (Taken from Figure 4 of Reference [48]).




FIG. 7. ESPS (‘Lattice Switch’) studies of the relative stability of *fcc* and *hcp* phases of the LJ solid at zero pressure, as discussed in section IV D. In this case the order parameter M (Eq. 59) measures the difference between the energy of a configuration of one phase and the corresponding configuration of the other phase generated by the switch operation. The areas under the two peaks reflect the relative configurational weights of the two phases. The evolution with increasing temperature (from *hcp*-favored to *fcc*-favored behavior) picks out the *hcp-fcc* phase boundary shown in Fig. 8. (Taken from Figure 7 of Reference [57]).




FIG. 8. A variety of approximations to the classical Lennard Jones phase diagram. The data points show the results of ESPS studies (discussed in section IV D), denoted here by ‘LS’. The dashed and solid lines are the results of harmonic calculations (for the two system sizes). The dash-dotted line is a phenomenological parameterisation of the anharmonic effects. The scale at the top of the figure shows the pressures at selected points on the (LS $N = 12^3$, NPT) coexistence curve. Tie-line structure is unresolvable on the scale of the figure. (Taken from Figure 11 of Reference [57]).




FIG. 9. The distribution of the density of the system of $N = 256$ LJ particles in crystalline and liquid phases, as determined by ESPS methods. The three pressures are (a) just below, (b) at and (c) just above coexistence for this N . (Taken from Figure 1 of Reference [58]).




FIG. 10. The distribution of the density of a LJ fluid at its critical point showing the collapse (given a suitable choice of scale) onto a form characteristic of the Ising universality class. (Taken from Figure 3a of Reference [44]).




FIG. 11. The difference between the free energy densities of *fcc* and *bcc* phases of particles interacting through a Yukawa potential, as a function of temperature, determined through the FG methods discussed in section VC. The error bars reflect the difference between the upper and lower bounds provided by FG-switches between the phases (along the Bain-path [85]) in the two directions. The favored phase changes around $T^* = 0.21$. (Taken from Figure 11 of Reference [84]).

ρ/ρ_{cp}	N	Δs ($10^{-5} \times k_B$)		Method	Ref.
0.736	12000	230	(100)	NIRM	[36]
0.736	12096	87	(20)	NIRM	[35]
0.7778	216	132	(4)	ESPS	[34]
0.7778	1728	112	(4)	ESPS	[34]
0.7778	1728	113	(4)	NIRM	[34]
0.7778	216	133	(3)	ESPS	[48]
0.7778	1728	113	(3)	ESPS	[48]
0.7778	5832	110	(3)	ESPS	[48]
1.00	12000	260	(100)	NIRM	[36]
1.0	216	131	(3)	ESPS	[48]
1.0	1728	125	(3)	ESPS	[48]

TABLE I. The difference in the *entropy* per particle of the *fcc* and *hcp* crystalline phases of hard spheres; the associated uncertainties are in parenthesis. For comparison we note that the excess entropy per particle at $\rho/\rho_{cp} = 0.7778$, $N = 1152$ is $-6.53 \dots$, with the phase-dependence showing in the 4th significant figure [28].

-
- [1] R. Evans in *Fundamentals of inhomogeneous fluids*, D. Henderson ed. (Dekker 1992).
 - [2] C. Caccamo, *Phys. Reports* **274**, 1 (1996).
 - [3] R.A. Cowley, *Advances in Physics* **12** 421 (1963).
 - [4] H.N.V. Temperley, J.S. Rowlinson and G.S. Rushbrooke, *Physics of simple liquids* (Amsterdam, North Holland, 1968).
 - [5] N.B. Wilding (unpublished).
 - [6] In the arguments of exponentials we shall generally absorb factors of $kT \equiv 1/\beta$ into effective energy functions.
 - [7] K. Binder and D.W. Heermann, *Monte Carlo simulation in statistical physics* (Springer, 1998).
 - [8] D.P. Landau and K. Binder, *A Guide to Monte Carlo Simulations in Statistical Physics* (Cambridge University Press, 2000).
 - [9] O. Narayan and A.P. Young, *Phys Rev E* **64** 021104 (2001).
 - [10] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller and E. Teller *J. Chem. Phys.* **21** 1087 (1953).
 - [11] It would be more circumspect to say that the partition function cannot in general be *naturally or usefully* expressed as a canonical average since that average inevitably falls into the category for which canonical sampling is inadequate [13].
 - [12] L.D. Fosdick, *Methods Comput. Phys.* **1** 245 (1963).
 - [13] In general the expectation value of an observable which is an *exponential* function of some *extensive* (‘macroscopic’) property cannot be reliably estimated from its sample average if one chooses $P_S = P_0$.
 - [14] Throughout this section we use the notation introduced in Eq. 7 in which the phase label is absorbed into an extended macrostate label C . The path defined by the set of macrostates $\{C\}$ may (and sometimes will) extend from one phase to another.
 - [15] We shall refer to a field λ and a macrovariable M as conjugate if the λ -dependence of the configurational energy is of the form $\mathcal{E}(\{q\}, c, \lambda) = \mathcal{E}(\{q\}, c) - \lambda M(\{q\})$
 - [16] It would be more conventional to refer to ‘densities’, which are *intensive*. But we prefer to deal with their *extensive* counterparts, which we shall refer to as *macrovariables*. The preference reflects our focus on simulation studies which are necessarily conducted on systems of some given (finite) size.
 - [17] Y. Iba, *Int. J. Mod. Phys. C* **12** 623 (2001).
 - [18] C.J. Geyer and E.A. Thompson, *J. R. Statist. Soc. B* **54** 657 (1992).
 - [19] K. Hukushima, H. Takayama and K. Nemoto, *Int. J. Mod. Phys. C* **3** 337 (1996).
 - [20] E. Marinari, *Optimized Monte Carlo Methods in Advances in Computer Simulation* ed J. Kertesz and I. Kondor (Springer Verlag, 1997).
 - [21] More explicitly $\mathcal{C}_1 \equiv \alpha, c$ and $\mathcal{C}_2 \equiv \tilde{\alpha}, c$.
 - [22] G.M. Torrie and J.P. Valleau, *Chem. Phys. Lett.* **28** 578 (1974); *ibid J. Comp. Phys* **23** 187 (1977).
 - [23] A.P. Lyubartsev, A.A. Martsinovski, S.V. Shevkunov and P.N. Vorontsov-Velyaminov, *J. Chem. Phys.* **96** 1776 (1992).
 - [24] E. Marinari and G. Parisi, *Europhysics Lett.* **19** 451 (1992).
 - [25] J.P. Valleau, *Adv. Chem. Phys.* **105** 369 (1999).
 - [26] M. Mezei M, *J. Comp. Phys.* **68** 237 (1987).
 - [27] B.A. Berg and T. Neuhaus *Phys. Lett. B* **267** 249 (1991); *ibid Phys. Rev. Lett.* **68** 9 (1992).
 - [28] D. Frenkel and A.J.C. Ladd, *J. Chem. Phys.* **81** 3188 (1984).
 - [29] D. Frenkel and B. Smit, *Understanding Molecular simulation* (Academic Press, 1996).
 - [30] J.P. Hansen and L. Verlet, *Phys. Rev.* **184** 151 (1969).
 - [31] D. Frenkel, H.N.W. Lekkerkerker and A. Stoobants, *Nature* **332** 822 (1988).
 - [32] W.G. Hoover and F.H. Ree *J. Chem. Phys.* **49** 3609 (1968).
 - [33] H Ogura, H. Matsuda, T. Ogawa, N. Ogita and A Ueda, *Prog. Theor. Phys* **58** 419 (1977).
 - [34] S. Pronk & D. Frenkel, *J.Chem. Phys.* **110** 4589 (1999).
 - [35] P.G. Bolhuis, D. Frenkel, S.-C. Mau and D.A. Huse, *Nature* **388** 235 (1997).
 - [36] L.V. Woodcock, *Nature* **384** 141 (1997); **388** 236 (1997).
 - [37] For hard spheres the Helmholtz free energy is purely entropic. Thus Table I shows the difference between the entropy per particle of the two phases, $[S_{fcc} - S_{hcp}]/N$.
 - [38] Q. Yan and J.J. de Pablo, *J. Chem. Phys.* **111** 9509 (1999).
 - [39] Equation 10 shows that the *logarithm* of the probability ratio measures the free-energy difference, which is *extensive*.
 - [40] In fact HPT seems to require that *all* the points chosen for replicas lie sufficiently close to coexistence that both phases have detectable signatures.
 - [41] B.A. Berg, U.H.E. Hansmann and T. Neuhaus, *Z. Phys. B* **90** 229 (1993).
 - [42] B. Grossmann, M.L. Laursen, T. Trappenberg and U.J. Wiese, *Phys. Lett. B* **293** 175 (1992).
 - [43] Reference [44] and N.B. Wilding (unpublished).
 - [44] N.B. Wilding, *Phys. Rev. E* **52** 602 (1995).

- [45] Some of the results from [44] are included in the LJ phase diagram shown in Fig. 4.
- [46] Operationally we may regard a phase as ‘metastable’ and assign it a ‘free energy’ (in relation to some other, also metastable) if the conditions for single phase sampling (see Eq. 21 and accompanying discussion) are met.
- [47] The configurations unearthed by the ES algorithm used in Lattice Switch studies of hard spheres (Section IV D 2) have some distinctive features [48].
- [48] A.D. Bruce, A.N. Jackson and G.J. Ackland, *Phys. Rev. E* **61**, 906 (2000).
- [49] See for example M.S.S. Challa and J.H. Hetherington, *Phys Rev A* **38** 6324 (1988).
- [50] When the inter-phase traverse proceeds by way of such two-phase configurations the free energy barrier to be surmounted must scale with the area of the interface supporting two phases in equal measure. The probability of the ravine-macrostates (and thus the probability of tunnelling between the two phases *at coexistence*) must vanish as e^{-F_I} where the free energy cost of the interface is of order $F_I \sim \sigma_s L^{d-1}$, with σ_s the interface tension and $L \sim N^{1/d}$ the edge of the d -dimensional simulation cube.
- [51] See for example reference [41].
- [52] The method to be discussed in the next section belongs in this family; so does HPT, but less compellingly because of the limitations we have already touched on in Section IV B 2.
- [53] See reference [29] p176.
- [54] The traverse might work acceptably well if the dynamics of the interface between the two phases is favorable: systems with martensitic phase transitions may fall into this category: Z. Nishiyama, *Martensitic Transformations* (Academic, New York, 1978). Note also that the special case in which the structural phase transition involves no change of symmetry can be handled within the standard multiccanonical framework [55].
- [55] G.R. Smith & A.D. Bruce, *Phys. Rev. E* **53** 6530 (1996).
- [56] A.D. Bruce, N.B. Wilding and G.J. Ackland, *Phys. Rev. Lett.* **79**, 3002 (1997).
- [57] A. D. Bruce, A. N. Jackson, G. J. Ackland, and N. B. Wilding, *Phys. Rev. E* **65** 36710 (2002).
- [58] N.B. Wilding and A.D. Bruce, *Phys. Rev. Lett.* **85** 5138 (2000).
- [59] In general the partition function Eq. 2 should already incorporate a Jacobean reflecting the transformation from cartesian coordinates to the generalized coordinates $\{q\}$. It may always be absorbed into the configurational energy and we have adopted that convention.
- [60] The constraint is rarely made explicit in MC studies: it is tacitly assumed that the conditions for ‘single-phase sampling’ (Eq. 21 and accompanying discussion) are met.
- [61] C.H. Bennett, *J. Comp. Phys.* **22** 245 (1976).
- [62] R. W. Zwanzig, *J. Chem. Phys.* **22** 1420 (1954).
- [63] A. Rahman and G. Jacucci, *Il Nuovo Cimento D* **4** 357 (1984).
- [64] One sees this more clearly when one views this relationship as the result of integrating Eq 55.
- [65] But a judicious choice of the *representations* chosen for the two phases, reflected in the transformation τ (Eq. 44) can bring one closer to this attractive limit. See references [70] and [71] and further remarks below.
- [66] We use the term ‘lattice’ rather loosely. The set of vectors $\{\vec{R}\}_\alpha$ identifies the orthodox lattice convoluted with the orthodox basis.
- [67] This will frequently be true for solid-solid phase behavior too; in the systems studied with Lattice Switch methods to date the effects have proved rather small.
- [68] M.C. Moody, J. R. Ray and A. Rahman, *J. Chem. Phys.* **84** 1795 (1985).
- [69] A.F. Voter, *J. Chem. Phys.* **82** 1890 (1985).
- [70] C. Jarzynski, *Phys. Rev. E* **65** 046122 (2002).
- [71] A.R. Acharya, A.D. Bruce and G. J. Ackland, in preparation.
- [72] A.Z. Panagiotopoulos, *Mol. Phys.* **61** 813 (1987).
- [73] A.Z. Panagiotopoulos, *J. Phys. Cond. Matt.* **12** R25, (2000).
- [74] J.P. Valleau *J. Chem. Phys.* **108** 2962 (1998).
- [75] K.K. Mon and K. Binder *J. Chem. Phys.* **96** 6989 (1992).
- [76] A.D. Bruce, *Phys. Rev. E* **55** 2315 (1997).
- [77] D. Möller and J. Fischer, *Mol. Phys.* **69**, 463 (1990); *ibid* **75** 1461 (1992).
- [78] B. Widom, *J. Chem. Phys.* **39** 2808 (1963).
- [79] A. Lotfi, J. Vrabec and J. Fischer, *Mol. Phys.* **76** 1319 (1992).
- [80] J. Vrabec and J. Fischer, *Mol. Phys.* **85** 781 (1995).
- [81] See for example H. Falk, *Am. J. Phys.* **38** 858 (1970).
- [82] C. Jarzynski, *Phys. Rev. Lett.* **78** 2690 (1997).
- [83] W.P. Reinhardt and J.E. Hunter III, *J. Chem. Phys.* **97** 1599 (1992).
- [84] M.A. Miller and W. P. Reinhardt, *J. Chem. Phys.* **113** 7035 (2000).
- [85] A.G. Khachatryan, *Theory of Structural Phase Transformations in Solids* (Wiley, NY, 1983).
- [86] A *continuous* space transformation like this can also be implemented in the lattice switch framework described in section IV D. By the same token one can generally devise a continuous version of any LS transformation. Thus, for example, the obvious continuous counterpart of the *fcc-hcp* LS transformation depicted in figure 6 entails a *gradual shear* of the

close-packed planes. Since the LS implementation remains always in the space of one or other of the crystalline structures one might expect it to prove the safer choice. Where an explicit comparison of the two strategies has been made [87] this was indeed the conclusion.

- [87] S.-C. Mau and D.A. Huse, *Phys. Rev. E* **59** 4396 (1999).
- [88] We use $\{\lambda\}$ to denote a set of physically different fields $\lambda_1, \lambda_2 \dots$. In contrast $[\lambda]$ (appearing in our discussion of paths) is shorthand for a set of different values $\lambda^{(1)}, \lambda^{(2)} \dots$ of some (particular) field λ . We use $\{M\}$ and $[M]$ in a similar way.
- [89] It reduces to the familiar form with the macrovariables M_1 and M_2 chosen as respectively the enthalpy and the volume whose conjugate fields are $1/T$ and P/T .
- [90] D.A. Kofke, *J. Chem. Phys.* **98** 4149 (1993); R. Agrawal and D. A. Kofke, *Phys. Rev. Lett.* **74** 122 (1995).
- [91] P.G. Bolhuis and D.A. Kofke, *Phys. Rev. E* **54** 634 (1996);
D.A. Kofke and P.G. Bolhuis, *Phys. Rev. E* **59** 618 (1999);
M. Lisal and V. Vacek, *Molecular Simulation* **18** 75 (1996);
F.A. Escobedo and J.J. de Pablo *J. Chem. Phys.* **106** 2911 (1997).
- [92] Recall that knowing the weights that define a multicanonical distribution of some quantity M over some range is equivalent to knowing the true canonical M -distribution *throughout the chosen range*.
- [93] For a review of FSS see *Finite Size Scaling and Numerical Simulation of Statistical Systems*, World Scientific Publishing, edited by V. Privman (Singapore, 1990).
- [94] C. Borgs and R. Kotecký, *J. Stat. Phys.* **61** 79 (1990); C. Borgs, R. Koteck and S. Miracle-Sol, *J. Stat. Phys.* **62** 529 (1991).
- [95] C. Borgs and W. Janke, *Phys. Rev. Lett.* **68** 1738 (1992).
- [96] As an immediate corollary, the field values that satisfy the equal-areas criterion will identify the coexistence curve of two such ensembles to within exponentially small corrections [95].
- [97] J.M. Polson, E. Trizac, S. Pronk and D. Frenkel, *J. Chem. Phys.* **112** 5339 (2000).
- [98] J.J. Binney, N.J. Dowrick, A.J. Fisher and M.E.J. Newman, *The Theory of Critical Phenomena: An Introduction to the Renormalization Group* (Oxford University Press, 1992).
- [99] See for example A.L. Sengers, R. Hocken and J. V. Sengers, *Physics Today* **30** 42 (1977).
- [100] A.D. Bruce and N.B. Wilding, *Phys. Rev. Lett.* **68** 193 (1992); N.B. Wilding and A.D. Bruce, *J. Phys.: Condens. Matter* **4** 3087 (1992).
- [101] M.R. Stapleton, D.J. Tildesley and N. Quirke, *J. Chem. Phys.* **92** 4456 (1990).
- [102] T. Kristof and J. Liszi, *Mol. Phys.* **99** 167 (2001).
- [103] P.V. Pant and D.N. Theodorou, *Macromolecules* **28** 7224 (1995).
- [104] N.B. Wilding and P. Sollich, *J. Chem. Phys.* **116** 7116 (2002).
- [105] L. Gomez, A. Dobry and H.T. Diep, *Phys. Rev. B* **63** 224103 (2001).
- [106] I.A. Rudnev, V.A. Kashurnikov, M.E. Gracheva and O.A. Nikitenko, *Physica C* **332** 383 (2000).
- [107] C.H. Bennett and B.J. Alder, *J. Chem. Phys.* **54** 4796 (1971).
- [108] S. Pronk and D. Frenkel, *J. Phys. Chem. B* **105** 6722 (2001).
- [109] D. Alfe, G.A. De Wijs, G. Kresse, M.J. Gillan *Int. J. Quant. Chem.* **77** 871 (2000).
- [110] Ecclesiastes 1:9. [The Holy Bible, King James Version (New York: American Bible Society, 1999).]
- [111] G.J. Ackland, *J. Phys. Condens. Mat.* **14** 2975 (2002).
- [112] F. Wang and D.P. Landau, *Phys. Rev. Lett.* **86** 2050 (2001); *Phys. Rev. E* **64** 056101 (2001).
- [113] F.H. Zong, D.M. Ceperley, *Phys. Rev. E* **58** 5123 (1998); C. Rickwardt, P. Nielaba, M.H. Muser and K. Binder, *Phys. Rev. B* **63** 045204 (2001).
- [114] M. Dijkstra, R. van Roij and R. Evans, *Phys. Rev. Lett.* **82** 117 (1999).
- [115] D. Frenkel, G.C.A.M. Mooij, B. Smit: *J. Phys. Cond. Matter* **4** 3053 (1992); S. Consta, N.B. Wilding, D. Frenkel and Z. Alexandrowicz, *J. Chem. Phys.* **110** 3220 (1999).
- [116] M. E. Fisher and A. Z. Panagiotopoulos, *Phys. Rev. Lett.* **88** 185701 (2002).
- [117] W. Kob, C. Brangian, T. Stuhn and R. Yamamoto, in *Computer Simulation Studies in Condensed Matter Physics XIII* eds. D.P. Landau, S.P. Lewis and H.B. Schtler (Springer, Berlin 2000).
- [118] M. Kolesik, M.A. Novotny and P.A. Rikvold, cond-mat 0207405.
- [119] In this section we refer to E rather than its dimensionless counterpart $\mathcal{E} \equiv \beta E$, since it seems advisable to make the measure of temperature $\beta = 1/kT$ explicit.
- [120] J. Lee, *Phys. Rev. Lett.* **71** 211 (1993); Erratum: *ibid* **71** 2353.
- [121] B.A. Berg *Int. J. Mod. Phys. C* **4** 249 (1993); *Fields Inst. Commun.* **26** 1 (2000).
- [122] G.R. Smith and A.D. Bruce, *J. Phys. A* **28** 6623 (1995).
- [123] Different authors have made different choices of notation and adopted different sign conventions at this point. We adopt those of [122].
- [124] The constant k merely serves to assign a convenient lower (or upper) bound to the weights; the value chosen is absorbed in the normalization of the sampling distribution.
- [125] The WL scheme is terminated when the histogram is ‘sufficiently’ flat *and* the weight-modification factor f is ‘sufficiently’ close to unity.

- [126] Q. Yan, R. Faller and J.J. de Pablo, *J. Chem. Phys.* **116** 8745 (2002).
- [127] There is a caveat here: the traverse through the interesting region needs to be slow enough to allow a rough local equilibrium in each macrostate to be established; this condition is satisfied relatively poorly at the outset (where the MC trajectory heads rapidly for the equilibrium states); but it gets better the closer the sampling distribution comes to the multicanonical limit; one requires no more than this.
- [128] P.M.C. de Oliveira, T.J.P. Penna and H.J. Herrmann, *Braz. J. Phys.* **26** 677 (1996); *ibid Eur. Phys. J. B* **1** 205 (1998); P.M.C. de Oliveira, *Eur. Phys. J. B* **6** 111 (1998).
- [129] J.S. Wang and L.W. Lee, *Comput. Phys. Commun.* **127** 131 (2000).
- [130] J.S. Wang, T. K. Tay, and R.H. Swendsen, *Phys. Rev. Lett.* **82** 476 (1999); J.S. Wang and R.H. Swendsen, cond-mat/0104418
- [131] The unspecified constant in Eq. B3 is prescribed by normalization.
- [132] We subsume into the definitions of the macrovariables any parameters which do not have the credentials ('linearity') of the fields $\{\lambda\}$.
- [133] A.M. Ferrenberg and R.H. Swendsen, *Phys. Rev. Lett.* **61** 2635 (1989); *ibid Phys. Rev. Lett.* **63** 1195 (1989); R. H. Swendsen, *Physica A* **194** 53 (1993).
- [134] The name seems to reflect no more than the fact that the resulting estimates of PDFs are as always based on (normalized) histograms accumulated in the sampling process.
- [135] The observable Q is then a delta-function that picks out prescribed values for the macrovariables in the set $\{M\}$.
- [136] To do so one must determine a set of weights that moderate the relative contributions of each histogram. There is more than a passing similarity here to the expanded ensemble framework discussed in Section III C.